

# Interpretivism and Inferentialism

David J. Chalmers

Robbie Williams' book *The Metaphysics of Representation* is the new leading edge of the program of naturalizing intentionality. Williams brings sophisticated ideas from many areas of philosophy to bear on this key problem in the philosophy of mind. Many philosophers will characterize their own views by how they relate to Williams' view. I'll do the same here.<sup>1</sup>

Williams' view is a form of interpretivism in the mould of Ramsey, Davidson, Lewis, and others. The basic idea of these approaches is that a subject is assigned the contentful mental states that best rationalize the subject's actions, given certain constraints and certain principles of rationality. The most familiar version of interpretivism appeals to means-end rationality as its central principle. Roughly, people have those beliefs and desires such that, if their beliefs are correct, their actions are likely to satisfy their desires.

This version of interpretivism comes up against a huge problem: underdetermination. Given one successful interpretation that maximizes means-end rationality, it is easy to permute the contents of beliefs and desires to come up with any number of alternative interpretations that serve means-end rationality equally well. To avoid this underdetermination problem, many have suggested adding further constraints: principles of charity, for example, or constraining contents to involve natural properties or properties that play certain causal roles. Still, underdetermination problems have often persisted even when these constraints are added, and no theory of this sort has met with any sort of wide acceptance.

In part to avoid these problems, Williams adds not one but three innovations to traditional interpretivism. First, he adds *source intentionality*. Williams' interpretivism presupposes a first tier of mental states in perception and action whose content is not grounded in interpretivism but instead through Karen Neander's teleosemantic analysis. Second, he appeals not just to structural rationality (means-end rationality, coherence, and the like) but also *substantive rationality*, including the rationality of induction, abduction, morality, and more. Third, he appeals to a *language of*

---

<sup>1</sup>Forthcoming in a symposium in *Analysis*.

*thought*, at least as a contingent feature of human thinkers, assigning contents initially not to total states of people but to word-like elements of structured mental states.

Each of these innovations is powerful and makes it much easier for interpretivism to meet various challenges. My worry is that the innovations are if anything too powerful, and that they may involve an element of overkill.

First: I think that if one appeals to source intentionality, one need not also appeal to substantive rationality to solve permutation problems.

Second: I think that if one appeals to a language of thought as Williams does, one moves (or should move) beyond interpretivism to a sort of inferentialism about the grounds of mental content.

As a result, I think Williams' powerful framework can be seen as supporting a sort of inferentialist metasemantics with source intentionality, or what I've elsewhere (Chalmers 2021) called a two-tiered inferentialism.

## **1 Source Intentionality and Substantive Rationality**

It's a familiar idea that taking some intentionality for granted can help to remove interpretivism's permutation problems. In Ramsey's pioneering form of interpretivism, he essentially took certain preferences for granted in choosing between possible outcomes. With preferences in hand, he was able to define relative utilities and credences, which are quantitative versions of beliefs and desires. Others have suggested taking the content of certain special words or representations such as *natural* or *cause* for granted, reducing or eliminating permutation problems that way.

Williams appeals to a distinctive sort of source intentionality: the intentionality of perception and action. You might think this would be more than enough to remove permutation problems. Surprisingly, Williams thinks his source intentionality is not enough to remove the problems. The reasons are tied to a permutation problem that Williams calls the "Bubble Puzzle". The basic idea is that the contents of perception and action always concern what is going on within a limited "bubble", so that one can always permute contents outside the bubble, yielding an unacceptable level of indeterminacy. As a result, Williams holds that in addition to source intentionality and structural rationality, one also needs to appeal to substantive rationality to remove permutation problems.

This raises obvious questions. We all know that appealing to enough source intentionality can remove permutation problems. Why appeal to source intentionality only to leave permutation

problems intact? Why not instead appeal to a slightly “thicker” version of source intentionality that removes these problems?

One simple way to go would be to appeal to source intentionality for certain basic concepts, corresponding to certain words in the language of thought. For example, it would be enough to appeal to source intentionality for spatial concepts, including concepts of spatial magnitudes such as being so-many meters apart. Suppose one’s perceptual bubble extends 100 meters from oneself. Then one can straightforwardly formulate thoughts such as “There are objects 200 meters away”: perhaps using different numerical magnitudes in thought, or perhaps by chaining representations of shorter distances such as one-meter-apart. The contents of these thoughts will be fixed by source intentionality, along with whatever it takes to fix logical structure. One can then have detailed permutation-resistant spatial beliefs about matters outside the bubble. With other perceptual concepts such as color predicates, one can fill in more details, and little of the original bubble puzzle remains. This way there’s no need for substantive rationality to avoid the bubble puzzle.

Perhaps Williams thinks it’s more acceptable to have source intentionality for perception than for cognition, but it’s not obvious why. It’s true that Neander’s teleosemantic story is presented as an account of the contents of perception, but it also seems straightforwardly applicable to perceptual concepts such as spatial concepts and color concepts. On the face of it, if Neander’s teleosemantic framework has the resources to assign contents to perceptual representations of redness or of distance based on their history, it should equally have the resources to assign contents to the conceptual representations involved in judgments such as *That is red* or *That is one meter away*.

If one doesn’t want to appeal to conceptual source intentionality, one can instead use perceptual source intentionality (which Williams already appeals to) to ground the content of perceptual concepts and thereby avoid the Bubble Puzzle. We need only assume that (1) perceptual representation is structured, with predicate-object structure and the like (e.g. attributing a certain shade of red and a certain location to a perceived object) (2) its predicative representations can attribute spatial properties and relations, such as the property of being one meter apart, and (3) structured representations in perception and cognition are connected in such a way that predicative representations in thought are quasi-inferentially/rationally connected to predicative representations in perception so that the former can inherit content from the latter. Versions of all of these assumptions are widely accepted by theorists of perception and cognition.

Things are particularly straightforward if we imagine that there is a language of perception as

well as a language of thought, and that spatial predicates can be shared between both. Then we can have a spatial relation predicate  $S$  where  $S(x, y)$  has the content that  $x$  and  $y$  are one meter apart.  $S$ 's occurrences in perception be assigned the content "one meter apart" by the principles of source intentionality. Its occurrences in thought will be constrained to have the same content (assuming we have a general constraint that occurrences of the same word in mental language have the same content). Then we can easily make determinate claims about the world outside the bubble as above.

Now, many deny that there is a common representational system for perception and for thought (e.g. Block forthcoming). But we don't really need a common system. Separate but connected structured representational systems are enough. It's plausible that perceptual states can attribute properties such as  $red_{31}$ , and that cognitive states can attribute the same properties. Furthermore, there's plausibly a relation of endorsement (or taking at face value) between perceptual and cognitive states, so that a perception with content *that is red<sub>31</sub>* is endorsed, yielding a judgment with the same content. If one doesn't think that perception and thought have exactly the same sort of content, there are various weaker relations that will do the same job, such as a class of endorsement transitions from perception to thought that are guaranteed to be truth-preserving. Then we our interpretivism need only invoke endorsement (given a perception, endorse it in thought) as a principle of structural rationality. Then Williams will be constrained to assign contents that respect this principle. In this way, spatial predicates in thought can inherit content from spatial predicates in perception, and we can repeat the process above to achieve determinate contents outside the bubble.

Williams might deny that endorsing perception in thought is a principle of rationality. Certainly it's a bit much to expect that we endorse everything we perceive; but many hold more weakly that when we do endorse a perceptual state in a belief, it's rational or justified to do so. Dogmatists and other "liberals" about perceptual justification typically hold a thesis like this. Many other "conservatives" don't, for example holding that perceptual states can't justify belief without justification for the additional premise that the perceptual state is accurate or reliable. Even a conservative might be able to find something rational nearby, however. For example, one could appeal to a transition from the  $red_{31}$  perception plus *Perception is reliable* or *This perceptual state is reliable* to the  $red_{31}$  belief. Or if there's such a thing as credences conditional on perceptual states, and one could argue that it's rational to have high credence in the  $red_{31}$  belief conditional on the  $red_{31}$  perception. One could also argue that imagination can be rational and that transitions from  $red_{31}$  beliefs to  $red_{31}$  imaginings count as rational. If so, a conservative interpretivist who

denies that the transition from  $red_{31}$  perceptions to  $red_{31}$  beliefs could still appeal to these other transitions to allow cognition to inherit content from perception.

These cases also raise questions about how strongly one's interpretivism should rely on constraints of rationality. In my view, what matters in the perception-belief case (and likewise the perception-imagination case) is not so much that the transition is rational as that it's truth-preserving and is naturally regarded as content-preserving. It's highly plausible that this sort of truth- and content-preserving transition provides constraints that a theory of content can exploit. If one's interpretivism can't use these constraints because it relies wholly on rationality and doesn't regard these inferences as rational, I'd say that's a strike against one's interpretivism. Perhaps one could take this as a point in favor of inferentialism (see the next section), which needn't rest on rationality and which can straightforwardly appeal to these transitions as content-constituting. But if one is committed to interpretivism, I'd strongly recommend a form that can allow these transitions to be content-constituting whether or not they count as strictly rational.

Now, Williams has a long discussion of observational concepts, with a complex story about how they refer to observational properties in virtue of the naturalness of these properties rationalizing our inductive practices. This story relies largely on substantive rationality, and as far as I can tell it hardly mentions perceptual source intentionality. Williams doesn't appear to consider a more straightforward story where where observational concepts inherit this content from perceptual states whose content is a matter of source intentionality. Perhaps this is because Williams thinks that perceptual states are unstructured or nonrepresentational, or perhaps it's because he thinks there are no rational transitions from perception to thought as above. Still, I think that current consensus strongly favors the view that perception is structured and representational at least in humans and other existing organisms. Neander's account of perceptual content, which Williams rests his account on, assumes the same view. Perhaps there are possible creatures in which perception is unstructured, but the same goes for thought (as Williams says), and it seems just as reasonable to appeal to the contingent architecture of structured representation for perception as for thought. Given this much, and given a view that appeals to perceptual source intentionality, it makes sense to allow that observational concepts can inherit content from perceptual states in some reasonably direct way as above.

I'm independently somewhat dubious about the appeal to substantive rationality—in particular the rationality of induction, abduction, and morality—as an interpretive constraint of agents. At this point I should confess that I don't really have a clear grasp of the distinction between structural and substantive rationality. But the rationality constraints I think are most relevant to

content individuation are those that turn on formal relations between contents, as found in deductive and means-end rationality, and possibly in perception-belief links (perhaps more generally in constraints tied to what I call “validity+” below). The rationality of induction, abduction, and rationality don’t seem as tightly tied to relations among contents, and it’s harder to see how they play a key role in individuating contents.

The difference can be brought out in imagining or interpreting specific agents. We have very little trouble imagining entirely amoral agents and even in interpreting them when confronted with them. Part of what makes them easy to imagine is that they may have quite coherent relations among their beliefs and desires, even if they believe and desire in irrational ways. We have only a little more trouble imagining counterinductive and counterabductive agents. These agents may behave oddly in not learning from experience, but there’s a straightforward behavioral profile for this. This contrasts strongly with imagining agents who are counterdeductive even for simple inferences such as modus ponens, or who are counterinstrumental for means-end rationality. It’s much harder to get a clear picture of what such an agent would be like. I take this as evidence that deduction and instrumental rationality are much more important than induction, abduction, and morality in our thinking about agents. I’m skeptical that induction, abduction, and morality play any constitutive role, though it would take more work to argue for this properly.

In any case, the complaint of overkill remains. From my perspective, Williams appeals to a huge resource of perceptual source intentionality, and then hardly uses the resource at key points where you might think it would be crucial, resting on substantive rationality instead. I’d say that theoretical economy strongly favors using one resource or the other but not both, if that’s possible. Perhaps it’s possible to use substantive rationality alone to eliminate underdetermination problems, though I’m not sure. I’m confident that one can use source intentionality alone to eliminate these problems, either through perceptual concepts or through perception itself, perhaps given minor plausible empirical assumptions and a tweak to one’s interpretism. I think there’s a strong case for trying out these more economical metasemantic views before bringing in both resources at once.

## **2 Interpretivism vs. Inferentialism**

The interpretivisms of Ramsey, Quine, Davidson, Dennett, and Lewis do not assume anything like a language of thought. They start with agents and their behavior, and interpret them from there. Williams rejects behaviorism for familiar reasons: Block’s giant look-up table might have all of our behavioral dispositions but none of our beliefs. To avoid this, we need constraints on internal

states. Williams does this by assuming structured representation in a language of thought, with reidentifiable symbols occurring in many different thoughts. To interpret agents as having mental states, we interpret these symbols as having whatever content maximizes the agent's rationality, in light of the constraints of source intentionality.

A language of thought is a vast resource for the interpretivist. It's nevertheless a fair resource for them to appeal to, in that proponents of rival theories (causal, teleological, and inferential) typically appeal to a language of thought, grounding intentionality by grounding contents for mental symbols. But one might wonder whether this resource transforms interpretivism into something very different from what it started as. It allows interpretivism to go far beyond pinning down contents by relations to actions, instead pinning them down through relations to structured thoughts and other mental states. Allowing this may bring interpretivism much closer to some rival views.

The rival view I'll focus on is inferentialism (or conceptual-role semantics), according to which the content of mental states and mental symbols (perhaps outside a layer of source intentionality) is grounded in their inferential role. This sort of inferentialism should be distinguished from inferentialist semantics (advocated by Block, Brandom, Harman, and others), according to which meanings and contents are inferential roles. Instead it is inferentialist metasemantics, on which meanings and contents (and in particular truth-conditional contents) are grounded in inferential roles.

I favor inferentialist metasemantics myself, but this isn't my only reason for focusing on the view. Williams' own view is explicitly inferentialist at many junctures. He appeals to the standard inferentialist account of the logical connectives, for example. He argues that his interpretivism, considered as a global account of intentionality that applies to all thoughts by all thinkers, supports an inferentialist account of intentionality in many local domains, such as thoughts involving the logical connectives. He justifies interpretivism in these domains by arguing that global interpretivism leads to local inferentialism, at least in these cases. So there's at the very least a strong inferentialist flavor to Williams' view. One can reasonably wonder why we shouldn't simply be inferentialist across the board.

There are two different questions one can raise here. First, is Williams' view itself a form of inferentialism? Second, if it's not, is there a nearby inferentialist view that we should prefer?

We can start with the first question. Williams' interpretivist view is that mental symbols have whatever content maximizes the rationality of the subject. A common form of inferentialism is that mental symbols have whatever content maximizes the validity of certain key inferences that the symbol is involved in. For example, Peacocke (1992) endorses a view on which the semantic

values of a concept is determined as what makes certain key inferences truth-preserving. So there's clearly a common flavor between Williams' interpretivism and traditional inferentialism. The potential differences are (1) Williams appeals to rationality-maximization, while inferentialists more often appeal to validity-maximization, and (2) Williams appeals holistically to the rationality of a subject, while inferentialism appeals locally to the rationality or validity of key inferences.

At this juncture, one might ask: is Williams' view a sort of holistic rationality-maximizing inferentialism, somewhat different from the more familiar local validity-maximizing inferentialism, but nevertheless clearly a form of inferentialism? Of course it doesn't matter much what we call the view, but if the differences between Williams' view and standard forms of inferentialism are as small as this, this is worth noting. Regimenting the views like this may also make it more straightforward to assess which view we should prefer.

To consider the potential differences in more detail: the first is whether the account invokes rationality (often characterized in normative terms, perhaps concerning inferences the subject ought to make or states the subject ought to be in) or validity (often characterized in semantic terms, perhaps involving truth-preservation). Both interpretivism and inferentialism can in principle come in both rationality-maximizing and validity-maximizing forms. We've already noted the possibility of validity-maximizing interpretivism in the previous section. Rationality-maximizing inferentialism may be less common than validity-maximizing forms in the literature, but it's nevertheless an available form of inferentialism. For example, if the "determination theory" in Peacocke's inferential analysis (1992) were tweaked to appeal to rationality rather than validity, the result would be a form of rationality-maximizing inferentialism. So if the only difference between Williams and inferentialists is his appeal to rationality as opposed to validity, then while this might be an advance, it would not be much of a step away from inferentialism.

For what it's worth, whether one is interpretivist or inferentialist in one's metasemantics, I favor something closer to validity-maximization than to rationality-maximization. One reason is tied to the perception- and imagination-involving transitions discussed in the previous section, with core transitions that may not be count as rational but certainly count as truth-preserving. Another reason is that I'm skeptical (as before) that the sort of rationality found in induction, abduction, and morality plays any role. A third is the deeply normative character of rationality. For a naturalistic account, Williams (in chapter 5) ends up having to invoke not rationality but a non-normative account of "rationality-making features"—something that it's not clear exists. In effect, Williams' naturalization of intentionality seems to require a naturalization of rationality (at least once intentionality is fixed). That's a huge task in itself if it's possible at all. Alternatively, one can

invoke a non-normative notion that falls short of rationality itself. If one does this, something like validity seems by far the clearest notion. Perhaps one should expand beyond validity construed as truth-preservation to a somewhat more expansive notion of “validity+” that covers non-truth-evaluable states, such as means-end reasoning involving desires and intentions (which may be valid in virtue of their satisfaction-conditions rather than their truth-conditions). But this will still be closer to validity-maximization than rationality-maximization.

For a validity-maximizing approach, there remains a difficult issue of which inferences one aims to maximize the validity of. It would perhaps be odd to maximize validity even for ampliative inferences that aren't presented as truth-preserving in the first place. So it's standard to restrict the class of relevant inferences, relying on a class of core, constitutive, or compelling inferences in fixing content. But then one faces the question of how to non-circularly delineate this class without presupposing the content of the states involved. There may be some benefits for a holistic rationality-maximizing view that doesn't have to face this issue. Overall, I think the balance of considerations favors something closer to validity-maximization, but both sorts of view are worth considering.

The second potential difference is Williams' holistic appeal to the overall rationality of the subject as opposed to the local rationality or validity of certain key inferences. To minimize the distance between the two, we can focus on a somewhat holistic rationality-maximizing form of inferentialism, where the semantic value of any mental symbol is determined to be whatever content maximizes the overall rationality of all inferences involving that symbol. Now the potential differences are (i) the rationality of a subject may involve more than the rationality of their inferences, and (ii) even if the rationality of a subject just involves rationality of inferences, it may be that maximizing the former holistically for all mental symbols at once gives different results from maximizing the latter locally for every mental symbol.

Regarding (i): Williams says explicitly that the rationality of an agent is determined by the presence or absence of rationality-making features of their beliefs and desires. If we assumed that (a) all such rationality-making features were features of the way those beliefs and desires were formed (based on evidence, for example), and (b) all methods of belief-formation count as inference, then the difference in (i) would disappear. Now, many inferentialists are happy to count all rational methods of belief-formation (including formation based on perception or intuition) as inference in a broad sense, so (b) is not a large difference. There remains (a): perhaps there are considerations such as consistency or coherence of beliefs and desires that are essentially synchronic rather than diachronic and are therefore not just matters of inference. Still, one could form

somewhat attenuated forms of inferentialism—coherence-involving inferentialism, for example—in which these matters play a role. My own form of two-tiered inferentialism gives coherence a role in assigning contents. So it's not clear that (i) is a huge difference between Williams' view and an inferentialist view.

Regarding (ii): it seems clear that locally maximizing the rationality of inferences for each mental symbol should also holistically maximize the rationality of inferences for the agent, at least if the local-maximization assignment is a possible assignment for the agent as a whole. Perhaps in some cases a local-maximization assignment is not a possible assignment for the agent as a whole. For example, perhaps maximizing rationality for A-involving inferences requires assuming a certain interpretation for B, one different from the interpretation that maximizes rationality of B-involving inferences. Inferentialists often propose apparatus to handle cases like this (my own view assigns contents locally in a certain order, for example) but this apparatus may lead to other divergences from a fully holistic approach. So perhaps there is some room for holism to make a difference. Still, it is not clear that the difference is large. And it is far from clear why, when there are differences, we should prefer the fully holistic approach to grounding content to the more local inferentialist approach.

This leads us to the second main issue above: even if there are some differences between Williams' interpretivism and inferentialism, which should we prefer? I'll contrast Williams' view with the sort of inferentialism that I favor: a two-tiered version that assumes a layer of non-inferentialist source intentionality. A first tier of mental symbols have their content as a matter of source intentionality, or inherited from source intentionality as in the previous section. A second tier have their content in virtue of inferential connections to symbols in the first tier and to each other, essentially maximizing the validity (across all scenarios, in the two-dimensional semantic framework) of certain key inferences.

Beyond this, a comparison between the views may rest largely on how they handle various domains. We can compare the two views first for the three domains that Williams considers in most depth: the logical connectives (chapter 2), observational concepts (chapter 3), and moral concepts (chapter 4).

Williams' treatment of logical connectives is explicitly a local inferentialist treatment. He justifies his interpretivism by arguing that given plausible assumptions (especially a "localizing" assumption saying that a holistically rational agent will find the relevant local inferences primitively compelling), it will yield the familiar local inferentialist results. That's a point in favor of inferentialism: it guarantees the results directly, without requiring the extra assumptions.

Something similar goes for moral concepts. Here Williams' central example is moral wrongness, and the central work is done by rational transitions from judgments about moral wrongness to attitudes of blame. He articulates a thesis saying that when these judgment-blame transitions are in place, the judgments involve wrongness. This looks like a local inferentialist treatment of wrongness, albeit in the inclusive sense where transitions between all sorts of mental states (not just judgments) can count as inferential. An inclusive inferentialist like me can take on board Williams' treatment of this concept wholesale. Again, the shape of this account seems to me to support inferentialism.

As for observational concepts, I suspect that Williams' complex account involving inductive inference could be straightforwardly be cast in inferentialist terms. In any case, I've already argued for a more straightforward view where these concepts inherit content from perception. That view suggests a form of inferentialism where perception-belief endorsement transitions play the central role in constituting belief content. We've seen that this picture may be slightly harder for a rationality-maximizing inferentialism (or interpretivism) to accommodate, depending on the rational status of these transitions, but it's no problem for a validity-maximizing inferentialism.

So I'd say that across these cases and the considerations above, there's a strong case for preferring an inferentialism where the content of a mental symbol is constituted by the transitions it's involved in. Perhaps Williams thinks there are other domains where inferentialism doesn't work as well and interpretivism does better. If so, it would be interesting to see those cases laid out. As things stand, Williams doesn't really present even a sketch of a metasemantic account for all sorts of familiar concepts—concepts such as *cat* and *in*, for example. My own inferentialism grounds their content fairly straightforwardly in their inferential relations to first-tier concepts (what sort of objects one counts as a cat, for example, and which scenarios one counts as one entity being in another, as well as what one can infer from these claims). Williams says at one point that he's open to a two-dimensionalist account like mine, but that he doesn't want to presuppose it, so it may be that he doesn't want to rely on the "scrutability" assumptions that ground my variety of inferentialism here. Still, I strongly suspect that when a Williams-style treatment of *cat* or *in* is spelled out, it will look similar in flavor if not in all details, and it won't give reason to favor interpretivism over inferentialism. That said, I'd love to see the details.

### **3 Inferentialism and Interpretivism as Foundational Theories**

In responding to an earlier version of these comments, Williams suggested that while inferentialism may often be the correct local story for a given thinker and a given domain of discourse, interpretivism provides a more plausible foundational story that may apply to all thinkers and all domains. One reason is that inferentialism requires a language of thought while interpretivism does not. I am doubtful that there is a significant difference between the theories here. There are state-based versions of inferentialism where the content of a belief is given by that belief's role in inference: for example, it has the content that maximizes validity of key inferences involving the state that the subject is disposed to make. A two-tiered inferentialism may constrain these key inferences to connect the state to first-tier states whose content is already fixed. Multi-tiered and holistic versions of state-based inferentialism are also available. So inferentialism seems to be at least on a par with interpretivism in this respect.

It's also worth noting that Williams appears not to give a state-based foundational story that doesn't involve a language of thought. His central framing of a foundational story (pp. 11-13) involves interpretations "mapping the atomic concepts onto contents, and containing rules for fixing the content of complex concepts and whole states out of the contents of their parts". He implies at some points that this framing is optional and that he could give a state-based foundational story that doesn't require a language of thought (while entailing his language of thought story when one is present), but he doesn't give such a story. So it's hard to assess the generality of Williams-style interpretivism without a language of thought. I can roughly envisage how such a story would go (identify certain states as beliefs and desires without identifying structured contents, and assign them those contents that maximize rationality in light of other constraints), but it's far from obvious how this story will entail Williams' own language of thought story in cases where a language of thought is present. How is exactly is a state-based story interpretivism sensitive to facts about representational structure that it doesn't presuppose as part of its grounds for interpretation?

Assuming these issues can be dealt with, then inferentialism and interpretivism both seem to be available as global foundational stories for grounding intentionality. Still, almost all of the local stories that Williams tells seem to be inferentialist, and in his book I don't see any obvious local cases that interpretivism can handle and inferentialism cannot. Perhaps Williams can spell out such a case. Certainly I'd be interested to see the strongest challenges for inferentialism here. If any such local challenges for inferentialism can be overcome, I think we then have strong evidence

for inferentialism over interpretivism as a foundational theory of intentionality.<sup>2</sup>

## References

Block, N. forthcoming. *The Border Between Seeing and Thinking*. Oxford University Press.

Chalmers, D.J. 2021. Inferentialism, Australian-style. *Proceedings and Addresses of the American Philosophical Association*.

Peacocke, C. 1992. *A Study of Concepts*. Oxford University Press.

Williams, J.R.G. 2020. *The Metaphysics of Representation*. Oxford University Press.

---

<sup>2</sup>Thanks to Adam Pautz and Robbie Williams for comments.