

Philosophical Explorations

An International Journal for the Philosophy of Mind and Action



ISSN: 1386-9795 (Print) 1741-5918 (Online) Journal homepage: <https://www.tandfonline.com/loi/rpex20>

Much ado about nothing? Why going non-semantic is not merely semantics

Daniel D. Hutto & Erik Myin

To cite this article: Daniel D. Hutto & Erik Myin (2018) Much ado about nothing? Why going non-semantic is not merely semantics, *Philosophical Explorations*, 21:2, 187-203, DOI: [10.1080/13869795.2018.1478988](https://doi.org/10.1080/13869795.2018.1478988)

To link to this article: <https://doi.org/10.1080/13869795.2018.1478988>



Published online: 02 Jul 2018.



Submit your article to this journal [↗](#)



Article views: 169



View related articles [↗](#)



View Crossmark data [↗](#)

Much ado about nothing? Why going non-semantic is not merely semantics

Daniel D. Hutto^{a*} and Erik Myin^b

^a*Department of Philosophy, University of Wollongong, Wollongong, NSW, Australia;* ^b*Department of Philosophy, University of Antwerp, Antwerp, Belgium*

This paper argues that deciding on whether the cognitive sciences need a Representational Theory of Mind matters. Far from being merely semantic or inconsequential, the answer we give to the RTM-question makes a difference to how we conceive of minds. How we answer determines which theoretical framework the sciences of mind ought to embrace. The structure of this paper is as follows. Section 1 outlines Rowlands’s (2017) argument that the RTM-question is a bad question and that attempts to answer it, one way or another, have neither practical nor theoretical import. Rowlands concludes this because, on his analysis, there is no non-arbitrary fact of the matter about which properties something must possess in order to qualify as a mental representation. By way of reply, we admit that Rowlands’s analysis succeeds in revealing why attempts to answer the RTM-question simpliciter are pointless. Nevertheless, we show that if specific formulations of the RTM-question are stipulated, then it is possible, conduct substantive RTM debates that do not collapse into merely verbal disagreements. Combined, Sections 2 and 3 demonstrate how, by employing specifying stipulations, we can get around Rowlands’s arbitrariness challenge. Section 2 reveals why RTM, as canonically construed in terms of mental states exhibiting intensional (with-an-s) properties, has been deemed a valuable explanatory hypothesis in the cognitive sciences. Targeting the canonical notion of mental representations, Section 3 articulates a rival nonrepresentational hypothesis that, we propose, can do all the relevant explanatory work at much lower theoretical cost. Taken together, Sections 2 and 3 show what can be at stake in the RTM debate when it is framed by appeal to the canonical notion of mental representation and why engaging in it matters. Section 4 extends the argument for thinking that RTM debates matter. It provides reasons for thinking that, far from making no practical or theoretical difference to the sciences of the mind, deciding to abandon RTM would constitute a revolutionary conceptual shift in those sciences.

Keywords: mental content; radically enactive cognition; computational theory of mind; representational theory of mind; teleosemantics

The rise of cognitivism is often regarded as nothing short of a Kuhnian revolution, with cognitive representations playing the leading role in defining the new paradigm. I’ve suggested here that there is now something like a counter-revolution taking place, even though most of the participants don’t realize it. (Ramsey 2007, 233)

*Corresponding author. Email: ddhutto@uow.edu.au

Could we, should we eliminate intentionality from sciences of the mind? To keep, or not to keep intentionality, that is *the* question. Philosophers are prone to grand proclamations when speculating about the consequences of answering it one way or the other. Rosenberg (2014) predicts a complete purge of intentional notions from mature scientific psychology. By his lights, we will only have a true understanding of our behavior if we accept that

whatever the brain does, it doesn't operate on beliefs and wants, thoughts and hopes, fears and expectations, insofar as these are supposed to be states that "contain" sentences, and are "about" things, facts, events that are outside of the mind. (2014, 26)

Taking the diametrically opposing view, Fodor (1987) presages that the loss of intentional psychology, "would be, beyond comparison, the greatest intellectual catastrophe in the history of our species" (Fodor 1987, xii).¹

It is easy to regard the above statements as shining examples of hyperbole: exaggerations brought on by nothing more than a taste for rhetorical flair. Yet, if we tone down their purely dramatic elements, such statements appear to convey an important truth. After all, is it not a pretty high-stakes question whether or not the sciences of the mind ought to retain a commitment to intentionality? *Prima facie*, it certainly seems so. It certainly looks to be a crucial question, especially if, as is standardly supposed, it translates directly into the question of whether the cognitive sciences ought to commit to a psycho-semantic – aka, to a representational theory of mind, or RTM. Belief in an RTM is indelibly linked to the familiar assumption that the primary business of minds is to manipulate contents of some kind. For ease of reference, let's call this the RTM-question.

Focusing on the contemporary scene, this paper argues, against detractors, that answering the RTM-question matters. Far from being merely semantic or inconsequential, the answer we give to the RTM-question makes a difference to how we conceive of minds. How we answer determines which theoretical framework the sciences of mind ought to embrace.

In making this case, the structure of this paper is as follows. Section 1 outlines Rowlands's (2017) argument that the RTM-question is a bad question and that attempts to answer it, one way or another, have neither practical nor theoretical import. Rowlands concludes this because, on his analysis, there is no non-arbitrary fact of the matter about which properties something must possess in order to qualify as a mental representation. By way of reply, we admit that Rowlands's analysis succeeds in revealing why attempts to answer the RTM-question simpliciter are pointless. Nevertheless, we show that if specific formulations of the RTM-question are stipulated, then it is possible to conduct substantive RTM debates that do not collapse into merely verbal disagreements. Combined, Sections 2 and 3 demonstrate how, by employing specifying stipulations, we can get around Rowland's arbitrariness challenge. Section 2 reveals why RTM, as canonically construed in terms of mental states exhibiting intensional (with-an-s) properties, has been deemed a valuable explanatory hypothesis in the cognitive sciences. Targeting the canonical notion of mental representations, Section 3 articulates a rival nonrepresentational hypothesis that, we propose, can do all the relevant explanatory work at much lower theoretical cost. Taken together, Sections 2 and 3 show what can be at stake in the RTM debate when it is framed by appeal to the canonical notion of mental representation and why engaging in it matters. Section 4 extends the argument for thinking that RTM debates matter. It provides reasons for thinking that, far from making no practical or theoretical difference to the sciences of the mind, deciding to abandon RTM would constitute a revolutionary conceptual shift in those sciences.

1. Full of sound and fury

Not everyone agrees that much, if anything, hangs on attempting to settle the RTM-question. Echoing a worry raised by Haselager, de Groot, and van Rappard (2003), Rowlands (2017) provides reasons for doubting that there is any practical or theoretical value in trying to do so.

The problem, as Rowlands sees it, is that the RTM-question is bound up in complex ways with how one conceives of mental representation, and that, in turn, depends on how one conceives of the representational relation. On Rowlands's analysis, how we conceive of mental representation necessarily embeds some or other account of what constitutes the relation that holds between a mental representation and that which it represents. If so, this putative fact has an important consequence: since different accounts of the representational relation abound, there is not one but a plethora of different conceptions of what constitutes a mental representation.²

On Rowland's analysis, the very idea of a mental representation is inescapably complex and inherently dis-unified. He concludes that the notion of mental representation is irretrievably equivocal – irretrievably equivocal in a way that robs us of the hope of ever getting a useful answer to the RTM-question.

The root problem, as he sees it, is this: there can be no meaningful answers to the RTM-question because the cognitive sciences do not operate with a singular, proprietary notion of mental representation. Mental representational talk is polysemic, plastic, and open-ended: the notion of mental representation is used in many diverse, incompatible and yet seemingly equally legitimate ways by cognitive scientists.

Thus, for some, it is enough for something to be a mental representation that it reliably causally co-varies with some worldly feature. For others, anything that carries information about some worldly feature, in a way that goes beyond mere causal covariance, qualifies as a mental representation. Yet, others maintain that it suffices for something to be a mental representation if it has an etiological proper function in guiding action with respect to some possible state of affairs. Others still will say it is enough that something plays a systematic functional role with respect to other processes. And so on, and on. Moreover, it is, of course, possible to combine any of these and other criteria, insofar as they are logically compatible, in order to produce even more refined and stringent accounts of the profiles that must be satisfied if something is to qualify as a mental representation.

It is the openness and flexibility of the multifarious ways that cognitive scientists conceive of mental representations that explains why debates over the indispensability or otherwise of positing such entities in the sciences of the mind are mired in irredeemable ambiguity. Even more fundamentally, it appears we have no principled way of choosing between these options. Ultimately, it seems:

any profile will, in its own way, be controversial. Thinking that any such profile can license the claim that *r* is a representation (or that it is not) will require that one be willing to reject theories of representation that others adopt (and, in their eyes, have good arguments in support of this adoption). (Rowlands 2017, 4428)

For these combined reasons, Rowlands (2017) rejects the familiar strategy of first attempting to characterize what is required for something to be a mental representation in terms of a certain set of properties and then checking to see if anything with the said profile actually does relevant cognitive work. By his lights, the lack of a governing, proprietary concept of mental representation results in attempts to answer the RTM-question only inspiring “a series of more or less loosely connected debates scarred by ambiguities,

misunderstandings, equivocations and proponents talking past opponents” (Rowlands 2017, 4226).

Of course, unless terms are agreed and kept fixed – at least, for the sake of argument – it is impossible to have a debate on any topic or to make progress on any important question. Hence, if opponents cannot agree in advance, even conditionally, about what counts as a mental representation, then there will be no point in asking the RTM-question. There is then no point in trying to decide if mental representations ought to be regarded as indispensable posits of mature sciences of the mind. The dismal assessment Rowlands (2017) offers, however, goes far beyond making this banal statement. His deep worry is that there simply is no principled, proprietary profile to be had for what counts as a mental representation.³

Consequently, any possible answer to the RTM-question will always be hostage to arbitrary choices and stipulations about how the notion of mental representation is used. If so, disagreements with respect to the RTM-question will be rendered theoretically and practically pointless because they are merely verbal. What some call representations others will not: it’s all just semantics.

Consider a case that we take to be quite important. It might be agreed, on all sides, that some structure *r* possesses a particular property, or combination of properties, and that *r* plays a central role in cognition. For example, it might be agreed that *r* lawfully co-varies with and structurally resembles some aspect of the environment. Moreover, it might be agreed that *r* has the biological function of enabling organisms to complete certain cognitive tasks precisely in virtue of *r*’s special properties. The trouble is that even if this much is agreed, it is always possible to deny that *r* qualifies as a *bona fide* mental representation. In other words, one might accept that *r* has the specified profile and plays the specified role and yet deny that *r* is a mental representation.

For those persuaded by Rowlands’s (2017) attempted exposé, a different kind of strategy might seem attractive at this juncture. Why not adopt a wait-and-see approach of discovering what does important work in cognition, and simply calling whatever does the relevant work a mental representation, whatever properties it may have? Clearly, the situation Rowlands (2017) highlights cannot be remedied by adopting this less front-loaded approach. This is because it is just another way of revealing the RTM-question to be pointless. As Ramsey (2017) emphasizes, going this way completely neuters and degrades the notion of mental representation, making it a mere honorific label which carries no explanatory value or scientific interest.

2. A rose by any other name

Can Rowlands’s (2017) pessimistic conclusion be resisted some other way? Is it possible to specify, in a principled way, which features suffice for something to qualify as a *bona fide* mental representation in order to make assessing what he calls “representation claims” stable enough to be able to ask the RTM-question in a way that is meaningful and worthwhile?

Rowlands (2017) thinks that what is required, but evades our grasp, is a singular notion of mental representation – a notion of mental representation characterized by an agreed minimal property or set of properties. Arguably, the best candidate for such a notion is the generic notion of mental representation that is supplied in countless cognitive science textbooks. Rey (2015) gives a quite standard characterization of that canonical notion:

“Representation” has come to be used in contemporary philosophy and cognitive science as an umbrella term to include not only pictures and maps, but words, clauses, sentences, ideas,

concepts, indeed, *virtually anything that is a vehicle for intentionality* (i.e. anything that stands for, “means”, “refers to”, or “is about something”). (171, emphasis added)

Under this characterization, mental representations are taken to be essentially content-involving states of mind. As Rey (2015) makes clear, it is intentional properties, understood as semantically robust properties, that provide the mark of the representational. Accordingly, something is a mental representation only if it possesses representational content.

It is this familiar, generic notion of mental representation – one that defines the class in terms of content-bearing states – that many in the field take to be indispensable to the sciences of the mind. Rescorla regards the representational content thesis as “a core doctrine” (2016, 17) in the sciences of the mind. Elaborating on this theme, he writes:

Philosophers and scientists use the phrase “mental representation” in many different ways. The type of mental representation that concerns me involves representational content ... [such that] a mental state has a content that represents the world *as being a certain way*. We can ask whether the world is indeed that way. These states are semantically evaluable with respect to such properties as truth, accuracy, and fulfilment. (17, emphasis added)

Echoing this same sentiment, Neander also tells us that something only counts as a mental representation if it is “used to represent a ... target *as being a certain way* that it might or might not be” (2017, 35, emphasis added).

Importantly, to characterize mental representations as contentful states of mind that do special work in cognition is to assume that they have special properties. Crucially, they entail the existence of unique kinds of satisfaction conditions. That is to say, they are not just intentional (with-a-t), in that they are directed at targets, but intensional (with-an-s), in that they specify their targets under a descriptive mode of presentation.⁴

Intensionality (with-an-s) is taken to be a necessary feature of any mental state that has representational content.⁵ The reasoning is simple: intensionality is implied if a mental state is to represent their targets *as being a certain way*. And only if targets are represented as being a certain way is it possible to get things right or wrong. Misrepresentation can only occur if contentful representations, in the broadest sense, make claims of some sort – namely, if they “say” something about how things stand with the world.

Even if this much is agreed, there is still scope for further disagreement and hence semantic slippage about which exact properties the generic notion of mental representation commits us to. For example, Cummins and Roth (2010) think of cognitive maps as mental representations that – though they represent the world as being a certain way – do not represent in the same way that sentences or beliefs represent. Accordingly, for these authors, mental representations, “do not have propositional contents. They are not candidates for truth-conditional semantics” (Cummins and Roth 2010, 188; see also Crane 2009; Gauker 2010; Burge 2010). As the case of mental maps reveals, there is scope for construing the representational content of mental states, even on the assumption that they are intensional, in terms of correctness conditions other than truth-conditions – such as, for example, accuracy or veridicality conditions.

With this caveat in place, we clarify the relatively wide scope of the generic notion of mental representation – minimally, a mental representation must be a contentful state of mind that exhibits intensionality and which does distinctively representational work in cognition. Arguably, the idea that cognition involves mental representations with these properties is an official, foundational assumption of cognitive science and is one of the main pillars of cognitivism (see Hutto and Myin 2017, Ch. 2). On the face of it, with the appropriate

stipulations in place, the generic notion of mental representation can surely provide a stable, hard target when conducting debates about whether cognitive science needs an RTM.

At this juncture, we can expect Rowlands to issue his arbitrariness challenge with renewed force. If we answer the question about whether intensionality features in the best cognitive explanations, then what, in addition, is gained by asking if mental representations play that very role? On Rowlands's analysis, the trouble with the RTM-question is that it is treated as an additional substantive question, when it isn't. And the reason he thinks it is superfluous is because there is no non-arbitrary fact of the matter about which properties something must possess in order to qualify as a mental representation. Hence, so understood, the RTM-question is simply decorative – it always boils down to a mere arbitrary choice of labels.

Should we not be worried by the fact that the canonical, textbook conception of a mental representation is not mandatory (Rowlands 2017, 4226)? It is certainly true that no one working in the field is obliged to adopt it. Still, for all that, we can sensibly ask whether the properties that the generic notion of mental representation takes to define the class really do play an essential, indispensable part in cognitive explanations. Nothing in Rowlands's (2017) analysis rules out asking or answering that question in a meaningful way, once the relevant stipulations are in place. And, crucially, once such stipulations are in place and agreed, the RTM-question cannot be treated as an additional question. As long as we are clear about the commitments of the generic notion, we will know which properties it assumes mental representations to have. This makes it possible to frame the RTM-question in such a way that answering it can be a substantive and not merely semantic business.

Price (2013) distinguishes between the activities of responding to and keeping track of covariant information and contentfully describing the world in ways that can be correct or incorrect. To capture this difference, he proposes using different labels for what he takes to be two different notions of representation. Crucially, he tells us that:

these two notions of representation should properly be kept apart, not clumsily pushed together. It takes some effort to see that the two notions of representation might float free of one another, but I think it is an effort worth making . . . Once the distinction between these two notions of representation is on the table, it is open to us to regard the two notions as having different applications, for various theoretical purposes. (Price 2013, 37)

Theoretically speaking, it is unimportant that Price (2013) uses the general label “representation” to denote two quite different phenomena: representation-as-covariation and representation-as-content-involving. Practically speaking, given the important differences he highlights between these two sorts of phenomena, it might court less confusion to reserve the word “representation” for only one of these phenomena and not the other. But that hardly matters: *Roses by other names!*

The Price (2013) strategy suggests that the right way forward is to make our notions of mental representation more precise and to tackle RTM-questions in a piecemeal way as opposed to asking it simpliciter. Put otherwise, many RTM-questions can be posed with respect to particular RTM theories and proposals. If this is correct, it will be entirely possible to conduct worthwhile investigations with respect to specific representation claims. This need not descend into chaos and the gnashing of teeth so long as everyone keeps their eyes firmly on the relevant properties and avoids quibbling over labels.

3. All that glitters is not gold

Once we know what a certain theory or family of theories proposes – namely, which properties are being claimed to do specific kinds of cognitive work – it becomes possible to have meaningful disagreements about the character of cognition and to investigate, in a substantive way, which properties cognition actually has. If we are careful in setting the stage for our debates, we can ask, without ambiguity, whether mental representations, as canonically characterized by a certain profile and job description, really ought to figure in the best explanations of cognitive work. The fact that the answers to RTM-questions will be controversial and hotly debated does nothing to show that the questions themselves are undecidable or unimportant, or that the debate over RTM is really just semantics.

It is entirely possible to ask, and advance substantive arguments for and against different answers to such questions as: Are mental representations – understood as entities with contentful properties that play a representational role in cognition – necessary for explaining intelligent action? Do the best explanations offered by working cognitive scientists really depend on positing such entities? Neander (2017) asks precisely these questions and makes a case for answering in the positive.

She maintains that the explanatory hypotheses that are standard fare in the cognitive sciences are committed to explaining behavior in terms of content-bearing representational structures. Thus, to use her own headline example, the explanation offered for how we fix the direction of a sound is in large part a story of content-rich processes through which the auditory system manages to use information concerning the inter-aural time disparity to work out where the sound comes from.

As Neander (2017) emphasizes, the hypothesized processes are unlike the slow inferential steps we would have to make in working this out explicitly. Indeed, unlike the explicit inferences that we might make, the processes of cognitive systems are assumed not to be conceptually based: hence the inferences the auditory systems make are not assumed to manipulate concepts such as “sound” or “inter-aural time disparity”. Nevertheless, it is conjectured that cognitive systems use nonconceptual representations that have representational contents understood in the canonical sense. For this reason, Neander acknowledges that “a red flag would be raised if the relevant content ascriptions were not intensional (with an ‘s’)” (2017, 36).

To allay this concern, Neander (2017) considers a second example of a familiar cognitive explanation, focusing on an explanatory hypothesis discussed in McCloskey (2009). The hypothesis in question emerged from the study of perceptual deficits, but makes conjectures about the nature and character of normally functioning cognitive structures and processes.

The explanatory hypothesis attempts to account for the selective deficit of a subject, AH, who had difficulty in reliably locating visual targets using visual cues alone. In testing her capacities, it was discovered that AH had normal acuity with respect to the distance of visual targets but made errors with respect to the direction of such targets. Moreover, AH could determine if what was displayed was, say, a cross or a circle, but she was unable to determine with accuracy where such items were located.

In a bid to get at the root cause of AH’s particular pattern of errors, it is hypothesized the fault lies in the content of structures in her visual system and how that system mentally represents the locations of visual targets. Positing a contentful variety of mental representations, cognitive scientists such as McCloskey (2009) purport to explain why AH manages to detect the distance of targets accurately but is inaccurate when it comes to detecting their direction.

Generalizing, Neander claims that in cognitive explanations of this familiar sort, “a notion of malfunction-permitting function and a notion of error-permitting representation figure *centrally, explicitly, and together*” (2017, 34). For this reason, she takes it that such explanations are committed to genuine representational content.

To show that this conclusion is justified, Neander (2017) observes that the explanatory hypothesis under scrutiny assumes that the visual system employs a fourfold representational schema. Captured in a Cartesian coordinate system, “+” and “-” can be used to represent locations to the right and left of the origin, respectively, and “+” and “-” can be used to represent locations up and down from the origin, respectively. On this basis, Neander tells us that states of AH’s visual system

might be described as representing a visual target as at (say) $-5, +10$. This means that the target is represented (at some level of processing in AH) as five units to the left and ten up from her point of visual attention. (2017, 37)

In line with the explanatory hypothesis, she asks us to “suppose that AH’s visual system represents a visual target (for example, a circle on a screen) as at $-5, +10$ ” (Neander 2017, 37). She then observes that there are other ways we could refer to this location. For example, it could be also picked out by the illocution she supplies: “the location mentioned above for illustrative purposes”. And there are, of course, many other ways to denote it too. Yet, she claims, among all the possible descriptions we might give of the state of the visual system, we have reason to think that the first one – according to which it represents the target as being at $-5, +10$ – is, in her words, “more accurate” (Neander 2017, 37). From this fact, she concludes that, “content ascriptions made with respect to AH’s subdoxastic processes are intensional” (Neander 2017, 37).

It may be that the content ascriptions to states of perceptual systems are assumed by those who make them to describe the intensional contents of such systems. Yet even if that is the case, does it get us any closer to establishing that perceptual systems, in fact, operate with intensional representations? The illustrative example of AH’s deficit shows that a state of her visual system can be given various intensional descriptions and that, for certain explanatory purposes, we may have excellent reasons to prefer some over others. Yet that, in itself, does not secure what defenders of RTM, canonically construed, need to establish – namely, that we have good reason to think the states of perceptual systems *themselves exhibit* intensionality.

More needs to be done in order to show that the preferred description of the state of a perceptual system is a description of the intensional content of that perceptual system itself – for instance, that the visual system is descriptively representing the target *as* being at $-5, +10$. Returning to Neander’s example, for all that has been said about AH’s visual system, it might be targeting the visual features at the relevant locations described wholly extensionally. That possibility is perfectly compatible with certain intensional descriptions of that targeting being more useful than possible others, for specific explanatory purposes, in depicting what the system is doing.

It would be a blatant non-sequitur to move from the fact that the visual system can be described, intensionally, as representing a target in a certain way to the conclusion that the visual system represents a target in a certain intensional way. Confusion about how intensional and extensional contexts operate could, ironically enough, cause one to fail to see the logical error here. In any case, attempting to secure such a conclusion by such means would clearly beg the question against the nonrepresentationalist. Therefore, at best, in considering

the example of AH, Neander (2017) demonstrates that it is logically possible that there might be mental representations that exhibit intensionality, but she does nothing more.

Does anything independently motivate accepting RTM's explanatory hypothesis as canonically construed? Neander (2017) suggests that contentful states of mind do special work that allows us to understand a range of mental phenomena – thus, positing such states has a distinctive explanatory payoff. It seems that contentful mental representations are needed for all sorts of cognitive work: to play an epistemic role in informing beliefs; disposing us to judge certain things; guiding our actions; informing our conscious experiences; and enabling us to voluntarily control, to some extent, certain sub-personal processes (see Neander 2017, 35).

Is positing contentful mental representations really the only, or even the best, way to explain such psychological phenomena? It might seem so, given the way Neander (2017) lays out our explanatory options. She limits our choice to either positing an RTM or positing states of mind with “mere (i.e. natural-factive as opposed to intentional) informational content” (Neander 2017, 27). If these are our choices we have every reason to prefer RTM, for it hardly seems promising to try to account for the aforementioned phenomena by appeal to content-free forms of cognition that reduce to “*mere* information processing” (2017, 35, emphasis original). Arguably, this forced choice vastly underestimates the set of explanatory possibilities, masking RTM's true competition. To correct for this, in the next section, we describe what we take to be RTM's much more attractive explanatory rival.

4. A winter of discontent?

Is there is any alternative to Neander's choice of trying to explain the phenomena of interest to cognitive scientists either by positing states with mental representational content, as canonically construed, or positing “mere information processing” activity? Is there a middle way between these two extremes? We think so.

It seems possible to conceive of the intentional – a signature mark of the mental – as fundamentally contentless and non-intensional in character. Our Radically Enactive account of Cognition, REC, proposes going precisely this way – relinquishing the idea that representational content is an inherent feature of the intentional.

Focusing first on intentionality, REC holds that it is variegated – that it comes in a variety of forms and that in its more primitive form it lacks representational content, despite targeting certain worldly offerings. Looking at the larger history of attempts to explicate intentionality, it is possible to disentangle strands in Brentano's classical formulation that, we contend, are unhelpfully run together by many philosophers and which fuel their intuitions that there can be no intentionality without representational content. Yet, if we relax such intuitions, there is ample conceptual room for thinking that there can be forms of intentionality which are nonrepresentational in character (Hutto and Myin 2017, Ch. 5).⁶ We concur with Muller's assessment that we would benefit from a more “nuanced understanding of intentionality” (2014, 170).

Moreover, we hold that it is not only possible to give a naturalistic account of contentless intentionality, there are powerful considerations that push us in this direction anyway. As things stand, teleosemantic theories of representational content are the top choice for anyone who hopes to naturalize intentionality. Such theories seek to account for representational contents and the special norms they require using nothing more than an appeal to natural selection and evolutionary biology (Millikan 1984; 2005). Such theories of content are widely deemed by many to be “the most promising” (Artiga 2011, 181), and for some clearly “the best”, if not the “only”, real options for naturalists (Rosenberg 2013, 4).

We concur. Even so, REC recommends embracing only a teleosemantics that has minimal ambitions – namely, one that does not seek to provide a representational theory of content. REC appeals to teleological theories in order to explain how basic forms of cognition which target worldly offerings could have arisen naturally. It assumes that the non-accidental success of such cognitive activity is explained by certain correspondences holding between states of the organism and states of the environment. Yet it agrees with critics of classic teleosemantics – such as Burge (2010) and Rescorla (2012) – that there is no evident explanatory advantage or obvious justification for thinking that biologically basic forms of cognitive activity require states of mind with representational content understood as exhibiting intensionality.

Understood in the way we propose, teleosemantics yields a theory of Ur-intentionality; it does a job that differs from the job that many proponents of teleosemantics ask it to do. On our account, the machinery of teleosemantics can be successfully used to account for the most basic forms of targeted but non-intensional cognition – forms of cognition that, we contend, are the roots of cognition.

Why deviate from standard readings of teleosemantic theories? Notoriously, despite their many attractive features, classic teleosemantic theories encounter a serious problem when it comes to assigning determinate intensional contents to mental representations. For how, invoking the favorite example, is it possible to determine the content of the mental representations that are purportedly involved in a frog's tongue snapping behavior if we rely only on selectionist explanations? Is the frog representing the object of its interest as a fly, as a small moving black dot, as a nutritious object, or as food (Fodor 1990, 2008)?

Calling on Sober's selection of/selection for distinction, several theorists have tried to solve this indeterminacy problem by advancing different versions of the Target of Selection Argument, or TSA (e.g. Sterelny 1990; Agar 1993). Such theorists propose different counterfactual analyses that attempt to determine what, in the past, was *de facto* targeted that answered organismic needs often enough to enable the cognitive system to proliferate. Different analyses yield different results. On some analyses it is food that is targeted, *not black dots nor flies*, and so on. On other analyses, it is flies, *not black dots, nor food* and so on. On yet other analyses, it is black dots, *not flies, nor food*, and so on.

In a seminal article, Goode and Griffiths (1995) argue that these answers are complementary, not competing – they stand in a means-end hierarchy. Once this is understood it becomes clear how to answer the indeterminacy problem about what cognitive systems target, on the assumption that all of the properties in question – being a fly, being food, being a moving black dot, and so on – mattered causally to relevant selection processes. If so, the following composite list might be thought to capture the content of the frog's representation: “small, dark and moving, nutritious, fitness-enhancing, fly”.

We resist the invitation to make any such intensional ascription to the frog's perceptual system. Nothing is lost if we assume that the list picks out the state of affairs that frogs are supposed to target in extension. Such targeted cognitive activity would be sensitive to structural covariances in the environment because of having a particular selective history. This can be wholly accounted for without assuming that basic forms of cognition are contentful in an intensional sense. Biological normativity does not suffice for norms of truth or accuracy where these norms are understood in representational terms as requiring satisfaction conditions that exhibit intensionality. This does not preclude the possibility that such targeting capacities may be a necessary platform for acquiring contentful attitudes, but it does preclude that such targeting capacities suffice for having contentful attitudes.

When teleosemantics is understood in this way it avoids Fodor's (1990) complaint that appeal to biological function cannot fix intensional contents because selectionist

explanations are transparent and extensional. Thus, he famously draws attention to the fact that “Darwin doesn’t care how you describe the intentional object of frog snaps ... Darwin cares how many flies you eat, but not what description you eat them under” (Fodor 1990, 73).

These influential worries about teleosemantics vanish entirely if we think of basic intentionality as having specific targets but not as specifying those targets in intensional ways. Consider, after all, that what Fodor objected to “wasn’t the tension between Darwinism and theories that are intentional (with a ‘t’) but the tension between Darwinism and theories that are intensional (with an ‘s’)” (Fodor 2008, 1).

Thus, in rejecting mental representations, REC holds that the most basic forms of cognition lack intensional content; they do not represent their targets *as* being a certain way. It would be a mistake, however, to assume – as Hufendiek (2016) does – that this entails that contentless forms of cognition “do not have adequacy conditions” (Hufendiek 2016, 96). REC assumes that there are certain ways of responding to aspects of the environment that were etiologically selected for by evolutionary and other adaptive pressures.⁷ Accordingly, selective forces set standards which any instance of organismic responding to worldly offerings might fail to satisfy. Thus, even without bringing intensionality or representational content into the equation, it is possible that an organism’s actual responses will be misdirected, amiss or awry with respect to some worldly target: an organism might fail to respond as it ought to respond.

In the end, REC proposes a position that is akin to the one Sachs (2012) advocates. Sachs maintains that it is only with respect to special forms of non-basic cognition that it makes sense to talk of content and is appropriate to “distinguish between sense and reference” (Sachs 2012, 145).

In going extensional, is not REC’s account of basic cognition too thin and colorless to compete explanatorily with its representationalist rivals? Returning to Neander’s list, she thinks we need intensionality (with-an-s) to explain: how mental states play an epistemic role in informing beliefs; dispose us to judge certain things; guide our actions; and inform our conscious experiences. We deny that we need to posit mental states possessing intensionality (with-an-s) to best explain such phenomena.

For example, do we need to posit contentful representations to explain what informs our conscious experience? REC holds that phenomenality – another signature mark of the mental – is, just like intentionality, fundamentally contentless, and non-intensional in character. Hence, it does not regard representational content as an inherent feature of the phenomenal.

REC’s account of basic cognition leaves room for worldly offerings to be experienced under aspects – the things creatures engage with can look or feel as certain way. But, we hold, such phenomenally charged ways of experiencing things neither entail nor are best explained by appeal to intensional contents (See Hutto 2006; Hutto and Myin 2013, Ch. 8). That basic cognition might be intentional but not intensional is consistent with perceptual experiences having particular phenomenal characters. There are ways that the world is experienced even if experiencing is not intrinsically contentful.

The qualitative aspects of perceptual experiences – how things look and feel to us – can incline or prompt the formation of contentful beliefs and judgments even if the experiences themselves do not “say” anything about the state of the world.

Raleigh (2015) defends the possibility that there can be phenomenal presentation without intensional representation. He asks, “Why can’t properties be presented in experience without them being presented ‘as being instantiated’?” (Raleigh 2015, 1213). Thus:

a single phenomenal look can correctly be described using many different predicates: “looks F”, “looks G”, “looks H” etc. ... The shared manifest way of looking in itself does not specify any one particular possible situation for the actual environment to match. (Raleigh 2015, 1236)

In specifying the content of an intensional representation we want to rule out certain possibilities about what might be represented. Yet when it comes to characterizing the qualitative character of an experience we precisely “*don’t want* to have a narrowly determinate scenario chosen as the ‘objective purport’ of the phenomenology” (Raleigh 2015, 1225).

Nor, pace Neander, is it true that perceptual experience must have intensional content if we are to explain how perceiving could play an epistemic role in grounding beliefs and judgments. In this regard,

The way things look can, in some contexts, reasonably be taken to provide evidence that something is the case, which is not in fact the case. But that the phenomenal look provides misleading evidence that something is the case does not show that the look is a kind of *claim* or *message* that something is the case. (Raleigh 2015, 1226, emphasis original)

This being so, there is no need to posit that perceptual systems have contents in order to account for the epistemic or justificatory links between how things look and how we believe them to be. A belief that the water is cold because it feels cold might be justified, under normal conditions, due to facts about how our sensory systems have been shaped to respond in such conditions.

5. A plague on both your houses

What if it turns out that the primary work of cognition isn’t to mentally represent? What if mental representations play no part in our best explanations of cognition? To avoid any ambiguity about which possible uses of the notion of mental representation is in play let us restrict our inquiries, as above, to investigating the consequences of eliminating RTM only as canonically construed. With that caveat in place, *pace* Rowlands, it seems that how we answer the above question matters – practically and theoretically – to the sciences of the mind. Indeed, if we relinquished RTM it would seem to matter a great deal. For if the standard stories about the most recent cognitive revolution are to be believed, then such an adjustment to how we think about minds would constitute a conceptual change of revolutionary proportions in the cognitive sciences (Hutto and Myin 2013, 2017).

Recalling the remarks of Rosenberg and Fodor that open this paper, how could the loss of RTM fail to have revolutionary repercussions? Is not cognitive science, after all, “the *representational* science of mind” (Clapin 2002, 19, emphasis original)? If we go by the textbooks, RTM is a – if not *the* – foundation upon which cognitive science is based (see O’Brien and Opie 2009; Shea 2013). Simply put, if RTM is in fact a cornerstone assumption of orthodox cognitive science, how then could abandoning it fail to constitute a pivotal conceptual change?

Nevertheless, some are skeptical that abandoning RTM would have any revolutionary import. Gärtner and Clowes (2017) criticize the claims that we have made about REC on this score. They question whether REC’s “anti-representational program is revolutionary” (Gärtner and Clowes 2017, 59).

One concern raised about REC’s putative status in this regard is that it proposes only very slight, and quite subtle, adjustments to existing theories. After all, REC’s main proposal is to offer a contentless teleosemantics – one which borrows nearly all of its theoretical

machinery from teleosemantics (Hutto and Myin 2013, 2017). This being the case, in a great many ways REC agrees with its RTM rivals (Gärtner and Clowes 2017, 77). This fact leads Gärtner and Clowes to conclude that it is far from clear that REC “offers anything new. At least it is not clear what is new about what is on offer” (2017, 71).

Hopefully, the preceding section makes clear both what REC offers that is different, and also why its subtle theoretical adjustment can have revolutionary significance. When it comes to conceptual revolutions, it is not the size of an adjustment that matters. Even a very subtle one, if it targets the right theoretical joint, can be a complete game changer.

We agree with Gärtner and Clowes that “the real question is whether REC is merely replacing a few sets of concepts or the whole system” (2017, 71; see Hutto and Myin 2017, Ch. 2). We think the ramifications of the slight change REC proposes to the classic teleosemantic theories initiate a domino effect – one with rippling ramifications for how we conceive of minds across the board.

There is a deeper worry; so Gärtner and Clowes (2017) contend. It is that REC simply cannot win. The reasoning is as follows: if the arguments for REC prove true, they establish that RTM as canonically conceived never mattered in the first place. And, if RTM never mattered in the first place, then its removal makes no positive contribution to our conception of cognition. Ergo, REC makes no positive contribution to our conception of cognition.

The perceived problem is that even if we adopt REC’s recommendations and abandon RTM, REC cannot be the basis for a conceptual revolution because, at best, it makes only a purely negative contribution to our understanding. REC only takes, it does not give; if it does anything it only eliminates a mistaken understanding of what does real cognitive work. For these reasons, Gärtner and Clowes (2017) doubt that REC’s proposed conceptual clarifications could possibly drive cognitive science forward.

In other words, at best, REC removes a conceptual blockage or barrier – RTM – that prevents us from thinking correctly about the primary work of cognition and how it gets done. But how, even if REC proves successful, could such a change positively contribute to the sciences of the mind? To make such a contribution REC would have to “restructure the fundamental concepts upon which a particular science is based” (Gärtner and Clowes 2017, 73). A key assumption at work in their reasoning is that “Only replacing *essential, fundamental* concepts of theories makes a conceptual revolution possible” (Gärtner and Clowes 2017, 73, emphases added). And with respect to cognitive sciences they assume, in line with REC, that “‘content’ and/or ‘representation’ are not concepts that describe the underlying mechanism” (Gärtner and Clowes 2017, 74).⁸

In the end, we reach Rowlands’ original conclusion about RTM by another route. If we accept such reasoning, the RTM debate is revealed to have neither theoretical nor practical import. Debates between representationalists and nonrepresentationalists, especially when the argument focuses only on the canonical conception of mental representation, are utterly futile and pointless. There is a plague on both houses!

Gärtner and Clowes (2017) maintain that it is the Computational Theory of Mind, aka CTM, and CTM alone – shorn of RTM – that is, and always has been, the sole and true foundation for the cognitive sciences. This is to adopt a wholly non-semantic vision of CTM, one which can be based on a mechanistic theory of computation, as recommended by Piccinini (2008, 2015), Fresco (2010) and Miłkowski (2013, 2016). The price of going this way is that to the extent that there was an initial cognitive revolution at all, one must hold that RTM should never have been billed as playing any part in it. If that were true, then the elimination of RTM from the sciences of the mind would, by

implication, not constitute a conceptual revolution. We take this attempt to re-write the standard history of the recent sciences of mind to be rooted in a confusion about the relation between epistemology and metaphysics.⁹

In closing, it is worth highlighting two worries about the CTM-only proposal. Firstly, there are reasons to doubt that a non-semantic CTM can do the work required by a theory of cognition. If anything fits the bill, CTM sans RTM is the best candidate for being the sort of “mere information processing” vision of mind that Neander (2017) decries and rates as inadequate for doing the required work of a theory of cognition – e. g. explain what informs and guides contentful attitudes and conscious experiences. For even if it turns out that there are good reasons to abandon classic RTM, there are equally good reasons for thinking that we need more than a purely non-semantic CTM if we are to account for the world-directed character of cognition.

Secondly, CTM might not be able to survive the loss of RTM in any case. Gärtner and Clowes recognize that “enactivism changes essentially how we evaluate the role of information” (2017, 76). Yet they ask if such a reevaluation is “enough for a revolution?” (2017, 76). Some, such as Villalobos and Dewhurst (2017a, 2017b), argue that CTM and a nonrepresentational enactivism are compatible and can live together. However, the radical form of enactivism we defend rejects not only the notion of intensional content, but also the very idea of informational content and informational processing (see Hutto and Myin 2013, 2017; Hutto et al. *in press*). In this respect, it challenges not only RTM but CTM.

If both pillars of cognitivism, RTM and CTM, were to collapse, then traditional cognitive science would have no basis – it would have no legs to stand on. And should that come to pass, it would be hard to deny that we were witnessing a revolution.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

Daniel D. Hutto thanks the Australian Research Council Discovery Project “Minds in Skilled Performance” (DP170102987) and Erik Myin thanks the Research Foundation Flanders (FWO) Project “Getting Real about Words and Numbers” (G0C7315N), as well as the Research Council of the University of Antwerp DOCPRO Project “Perceiving Affordances”, for funding the primary research informing this article.

Notes

1. In a similar vein, Baker (1988) proposed that elimination of intentional psychology would be tantamount to a kind of cognitive suicide.
2. Rowlands (2017) illustrates,

For example, that an item stands in a specifiable informational relation to another might make it a representation from the point of view of informational approaches but is completely irrelevant from the point of view of teleosemantic alternatives. That a state of a mechanism possesses a certain etiology will, similarly, be irrelevant from an informational or causal perspective. (p. 4220)

3. Rowlands (2017) holds that trying to answer the RTM-question is pointless, “even if there is a fact of the matter concerning whether an item is, or is not, a representation” (4216). However,

it is difficult to imagine what it might mean for there to be a fact of the matter about something being a representation if the notion of mental representation is irretrievably ambiguous in the way Rowlands suggests.

4. As Neander (2017) reminds us, there is no requirement that the subject “grasp an intension for a content ascription to qualify as intensional” (37). This fact is especially important when it comes to thinking about the intensionality (with-an-s) of nonconceptual mental representations. Moreover, we should not rule out a priori the possibility of such mental representations simply because we cannot find evidence for them in the absence of patterns of linguistic assent and dissent. For, as Neander (2017) rightly notes such patterns are “only easily accessible evidence of opacity. They are not essential for it” (38).
5. Neander (2017) is explicit about this commitment:

If a mental state represents at all, it represents something in one way and not in innumerable other possible ways. A young child can think that dawn has come before she can think, opaquely speaking, that the Earth has turned on its axis far enough to reveal the vast ball of flaming gas around which it orbits. (38)

6. In maintaining that intentionality is not always and everywhere permeated with representational content REC does not deny that it can be. Thus, in proposing that basic forms of Ur-intentionality are contentless, REC only challenges those theories which assume that representational content is either a, or the, defining feature of intentionality. Thus, contra Searle (1983), REC holds that we should not assume that the properties of all varieties of intentional attitudes are the same as those of contentful linguistic utterances and speech acts.
7. As Gärtner and Clowes (2017) observe, “teleosemantics need not be interpreted only in terms of classical Darwinian evolution, but could be interpreted in terms of any adaptationist process that could give rise to complex design or organisation” (Gärtner and Clowes 2017, 66).
8. Or, as Gärtner and Clowes (2017) put it, “clarifying or radicalizing an already existing view about an underlying mechanism does not change the mechanism” (71).
9. To adopt the line of thought of Gärtner and Clowes (2017) across the board, for all scientific and philosophical concepts, would make it impossible to account for intellectual progress. Their picture of concepts misconstrues how conceptual progress in general is made in science, and philosophy. Acquiring concepts is an achievement, or a process where a stable end product is gained only through elimination, refinement and elaboration. The concept of inertia, for example, was not obtained in one step neither by Galileo or Newton. Galileo’s concept, constructed, with crucial innovations, from existing concepts, still contained errors (such as allowing for circular inertia), which only could be rectified in the theoretical framework Newton built (Dijksterhuis 1961).

Notes on contributors

Daniel D. Hutto is Professor of Philosophical Psychology at the University of Wollongong and member of the Australian Research Council College of Experts. He is co-author of the award-winning *Radicalizing Enactivism* (MIT, 2013) and its sequel, *Evolving Enactivism* (MIT, 2017). His other recent books include *Folk Psychological Narratives* (MIT, 2008) and *Wittgenstein and the End of Philosophy* (Palgrave, 2006). He is editor of *Narrative and Understanding Persons* (CUP, 2007) and *Narrative and Folk Psychology* (Imprint Academic, 2009). A special yearbook, *Radical Enactivism*, focusing on his philosophy of intentionality, phenomenology, and narrative, was published in 2006. He is regularly invited to speak not only at philosophy conferences but at expert meetings of anthropologists, clinical psychiatrists/therapists, educationalists, narratologists, neuroscientists, and psychologists.

Erik Myin is Professor of Philosophy at the University of Antwerp and director of the Centre for Philosophical Psychology. He has published papers on topics relating to mind and perception in philosophical, interdisciplinary and scientific journals. Two books, *Radicalizing Enactivism: Basic Minds without Content*, and *Evolving Enactivism: Basic Minds Meet Content*, written by Dan Hutto and Erik Myin, were published by MIT Press in 2013 and 2017.

References

- Agar, N. 1993. "What do Frogs Really Believe?" *Australasian Journal of Philosophy* 71 (1): 1–12.
- Artiga, M. 2011. "On Several Misuses of Sober's Selection for/Selection of Distinction." *Topoi* 30: 181–193.
- Baker, L. R. 1988. *Saving Belief: A Critique of Physicalism*. New Jersey: Princeton University Press.
- Burge, T. 2010. *The Origins of Objectivity*. New York: Oxford University Press.
- Clapin, H. 2002. "Introduction." In *Philosophy of Mental Representation*, edited by H. Clapin. New York: Clarendon Press.
- Crane, T. 2009. "Is Perception a Propositional Attitude?" *Philosophical Quarterly* 59: 452–469.
- Cummins, R., and M. Roth. 2010. "Meaning and Content in Cognitive Science." In *The World in the Head*, edited by R. Cummins, 174–193. Oxford: Oxford University Press.
- Dijksterhuis, E. 1961. *The Mechanization of the World Picture*. Oxford: Oxford University Press. Translation by C. Dikshoorn of E. Dijksterhuis, *De Mechanisering van het Wereldbeeld* (Meulenhoff, 1950).
- Fodor, J. A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. 1990. *A Theory of Content and Other Essays*. Cambridge: MIT Press.
- Fodor, J. A. 2008. "Against Darwinism." *Mind and Language* 23: 1–24.
- Fresco, N. 2010. "Explaining Computation Without Semantics: Keeping it Simple." *Minds and Machines* 20 (2): 165–181. doi:10.1007/s11023-010-9199-6.
- Gärtner, K., and R. Clowes. 2017. "Enactivism, Radical Enactivism and Predictive Processing: What is Radical in Cognitive Science?" *Kairos. Journal of Philosophy and Science* 18: 55–83.
- Gauker, C. 2010. *Words and Images. An Essay on the Origin of Ideas*. Oxford: Oxford University Press.
- Goode, R., and P. E. Griffiths. 1995. "The Misuse of Sober's Selection of/Selection for Distinction." *Biology and Philosophy* 10: 99–108.
- Haselager, P., A. de Groot, and H. van Rappard. 2003. "Representationalism vs. Anti-Representationalism: A Debate for the Sake of Appearance." *Philosophical Psychology* 16 (1): 5–24.
- Hufendiek, R. 2016. *Embodied Emotions: A Naturalist Approach to a Normative Phenomenon*. London: Routledge.
- Hutto, D. D. 2006. "Unprincipled Engagements: Emotional Experience, Expression and Response." In *Consciousness and Emotion: Special Issue on Radical Enactivism*, edited by R. Menary, 13–38. Amsterdam: John Benjamins.
- Hutto, D. D., and E. Myin. 2013. *Radicalizing Enactivism: Basic Minds Without Content*. Cambridge, MA: MIT Press.
- Hutto, D. D., and E. Myin. 2017. *Evolving Enactivism: Basic Minds Without Content*. Cambridge, MA: MIT Press.
- Hutto, D. D., E. Myin, A. Peeters, and F. Zahnoun. In press. "The Cognitive Basis of Computation: Putting Computation in Its Place." In *The Routledge Handbook of the Computational Mind*, edited by M. Colombo, and M. Spervak. London: Routledge.
- McCloskey, M. 2009. *Visual Reflections: A Perceptual Deficit and its Implications*. Oxford: Oxford University Press.
- Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories*. Cambridge: MIT Press.
- Millikan, R. G. 2005. *Language: A Biological Model*. New York: Oxford University Press.
- Milkowski, M. 2013. *Explaining the Computational Mind*. Cambridge, MA: MIT Press.
- Milkowski, M. 2016. "Function and Causal Relevance of Content." *New Ideas in Psychology* 40: 94–102.
- Muller, H. D. 2014. "Naturalism and Intentionality." In *Contemporary Philosophical Naturalism and Its Implications*, edited by B. Bashour and H. D. Muller, 155–181. London: Routledge.
- Neander, K. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. Cambridge: MIT Press.
- O'Brien, G., and J. Opie. 2009. "The Role of Representation in Computation." *Cognitive Processing* 10 (1): 53–62.
- Piccinini, G. 2008. "Computation Without Representation." *Philosophical Studies* 137 (2): 205–241.
- Piccinini, G. 2015. *Physical Computation: A Mechanistic Account*. Oxford: Oxford University Press.
- Price, H. 2013. *Expressivism, Pragmatism and Representationalism*. Cambridge: Cambridge University Press.

- Raleigh, T. 2015. "Phenomenology Without Representation." *European Journal of Philosophy* 23 (4): 1209–1237.
- Ramsey, W. 2017. "Must Cognition be Representational?" *Synthese* 194 (11): 4197–4214.
- Rescorla, M. 2012. "Millikan on Honeybee Navigation and Communication." In *Millikan and Her Critics*, edited by D. Ryder, J. Kingsbury, and K. Williford, 87–102. Oxford: Blackwell-Wiley.
- Rescorla, M. 2016. "Bayesian Sensorimotor Psychology." *Mind and Language* 31 (1): 3–36.
- Rey, G. 2015. "Representation." In *The Bloomsbury Companion to the Philosophy of Mind*, edited by J. Garvey, 171–189. London: Bloomsbury.
- Rosenberg, A. 2013. "How Jerry Fodor Slid Down the Slippery Slope to Anti-Darwinism, and how we can Avoid the Same Fate." *European Journal for Philosophy of Science* 3: 1–17.
- Rosenberg, A. 2014. "Disenchanted Naturalism." In *Contemporary Philosophical Naturalism and Its Implications*, edited by B. Bashour and H. D. Muller, 17–36. London: Routledge.
- Rowlands, M. 2017. "Arguing About Representation." *Synthese* 194 (11): 4215–4232.
- Sachs, C. B. 2012. "Resisting the Disenchantment of Nature: McDowell and the Question of Animal Minds." *Inquiry* 55 (2): 131–147.
- Searle, J. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Shea, N. 2013. "Naturalising Representational Content." *Philosophy Compass* 8 (5): 496–509.
- Sterelny, K. 1990. *The Representational Theory of the Mind*. Oxford: Blackwell.
- Villalobos, M., and J. Dewhurst. 2017a. "Why Post-Cognitivism Does not (Necessarily) Entail Anti-Computationalism." *Adaptive Behaviour* 25 (3): 117–128.
- Villalobos, M., and J. Dewhurst. 2017b. "Enactive Autonomy in Computational Systems." *Synthese*. doi:10.1007/s11229-017-1386-z.