

Neo-Pragmatism, Primitive Intentionality and Animal Minds

Laura Danón¹ 

Received: 21 September 2017 / Revised: 19 February 2018 / Accepted: 6 March 2018 /

Published online: 16 March 2018

© Springer Science+Business Media B.V., part of Springer Nature 2018

Abstract According to Hutto and Satne (*Philosophia*, 43(3), 521–536, 2015), an “essential tension” plagues contemporary neo-Pragmatist accounts of mental contents: their explanation of the emergence and constitution of intentional mental contents is circular. After identifying the problem, they also propose a solution: what neo-Pragmatists need to do, to overcome circularity, is to appeal to a primitive content-free variety of intentionality, different from the full-blown intentionality of propositional attitudes. In this paper, I will argue that, in addition to the problem of circularity, there is another important problem that both neo-Pragmatist accounts, and Hutto and Satne’s refinement of them, should also deal with: their difficulty to accommodate a host of recent empirical evidence and theoretical developments on the interdisciplinary field of animal cognition. I will call this difficulty *the objection from animal minds*, and I will present several arguments designed to show that, even though the notion of primitive intentionality, introduced by Hutto and Satne, may be useful to account for some of the most basic ways of dealing with the environment of nonhuman animals, it falls short of providing an adequate explanation of the full-range of cognitive capacities and behavioral dispositions that many animal species display. Thus, their proposal ends up being insufficient to help neo-Pragmatist approaches to overcome the problem of animal minds. Finally, I will suggest that overcoming this objection requires attributing to (several species of) non-human animals some basic, yet content involving, kinds of intentional mental states.

Keywords Primitive intentionality · Animal minds · Neo-pragmatism · Mental content

✉ Laura Danón
ldanon@unc.com.ar

¹ Centro de Investigaciones de Filosofía y Humanidades, Facultad de Filosofía y Humanidades, Universidad Nacional de Córdoba (National University of Córdoba), Pabellón Agustín Tosco s/n Piso 1°, Ciudad Universitaria, Córdoba, Argentina

In recent times, several philosophers have leaned towards what may be loosely called neo-Pragmatist accounts of mental contents (Sultanescu 2015; Cash 2008; Brandom 1994, 2000; Sachs 2014; Haugeland 1990). In a paper dedicated to critically assessing different contemporary attempts to naturalize mental contents, Hutto and Satne (2015) point out that an “essential tension” plagues such neo-Pragmatist accounts: i.e., their explanation of the emergence and constitution of intentional contents is a circular one. Yet, after identifying the problem, they also propose a solution: what neo-Pragmatists need to do, in order to overcome circularity, is to appeal to a primitive content-free variety of intentionality, different from the full-blown intentionality of propositional attitudes.

While I agree with Hutto and Satne’s claim that neo-Pragmatists must find a way to avoid the problem of circularity, I think there is another important problem that they should also deal with: their difficulty to accommodate a host of recent empirical evidence and theoretical developments in the interdisciplinary field of animal cognition. In this paper, I would like to evaluate whether Hutto and Satne’s proposal can indeed help neo-Pragmatist approaches to overcome this second kind of difficulty. As I will argue, even though their proposal may be useful to account for some of the most basic ways of dealing with the environment that non-human animals display, it falls short of providing an adequate explanation of the full range of cognitive capacities and behavioral dispositions that many animal species display. In what follows, I will call this difficulty *the objection from animal minds* and I will argue that, in order to overcome it, it is necessary to attribute to these creatures some kinds of mental contents.

Admittedly, the authors do not explicitly attempt to respond to the objection from animal minds. Nevertheless, if the objection as I will present it is correct, relevant consequences follow from it, both for the original neo-Pragmatist accounts and for Hutto and Satne’s attempt to amend them by adding a content-free variety of intentionality. Firstly, even though distinguishing between primitive content-free intentionality and full-blown intentionality may be an interesting theoretical move, it will not be sufficient to render such a strong version of neo-Pragmatism as the authors pretend (after all, the new version will still face difficulties to deal with evidence regarding the cognitive capacities of some non-human animals). Secondly, and more specifically, this insufficiency will be problematic for Hutto and Satne’s general attempt to provide a naturalistic account of the natural origins of intentionality. As I will explain later on, the authors are interested in elaborating a genealogical explanation of the origins of mental content in the natural world. But, if they can’t overcome the objection from animal minds, their proposal won’t be able to provide an empirically adequate account of the transition from entirely non-intentional states and processes to the complex and sophisticated varieties of human full-blown intentionality that neo-Pragmatists usually focus on. Thirdly, Hutto and Satne explicitly commit themselves to a methodological stance, which they dub as “Relaxed Naturalism”. According to it, philosophers should attempt to “clarify the nature of some *explanandum* by investigating it in a way that draws on and seeks to harmoniously integrate the findings of a wide range of empirical sciences” (p. 530). Now, if my objection is correct, there is a host of well-established empirical evidence from different scientific disciplines that their account will not be capable of integrating harmoniously, and this is, clearly, something that they, being Relaxed Naturalists, should worry about. It seems, then, that both neo-pragmatists accounts of content and Hutto and Satne’s proposal would benefit if they could find a way to overcome the problem of animal minds. However, as I will argue, to accomplish that aim it is necessary to posit not only primitive content-free

intentionality, but also some basic variety of content-involving intentionality. Only such an addition will allow explaining the kinds of complex and flexible behavioral patterns displayed by many non-human species.

Here is the structure of the paper. I will begin by reconstructing – admittedly in a brief and rough way—the main tenets of recent neo-Pragmatist approaches to intentionality, and I will spell out the objection of circularity that Hutto and Satne raise against them, as well as their proposal on how to enrich neo-pragmatists’ accounts of intentionality by positing a new variety of contentless intentionality (section 1). In section 2, I will present the objection from animal minds and I will offer several arguments designed to show why the notion of primitive intentionality introduced by Hutto and Satne is insufficient to solve this second problem. In section 3, I will briefly examine some possible ways in which a defender of primitive intentionality might try to respond to the objection from animal minds, by leaning on some recent developments by Radical Enactivists, like Hutto and Myin (2017). As I will argue, even if they do provide some promising ways of enriching the notion of primitive intentionality and extending its explanatory potential, these new theoretical proposals do not suffice to overcome the problem. Finally, I will suggest that what enables us to overcome objection from animal minds is attributing to (several species of) non-human animals some basic, yet content involving, kinds of intentional mental states.

1 Neo-Pragmatist Conceptions of Mental Content

Neo-Pragmatist approaches to intentionality –defended, amongst others, by philosophers such as Sellars, Davidson, Brandom, Sachs, Cash, and Haugeland— are committed to different versions of three basic claims. Firstly, intentional mental states are derived from – and can only emerge in – the context of normative socio-cultural practices (Haugeland 1990; Hutto and Satne 2015; Satne 2016; Sultanescu 2015). Secondly, in order to have mental states endowed with intentional contents, a creature needs to be actively engaged (or at least to be capable of participating) in normative social and cultural practices (Brandom 1994, 2000; Hutto and Satne 2015; Haugeland 1990). Thirdly, an adequate account of mental content must crucially and unavoidably rely on the notion of participation in shared practices (Sultanescu 2015).

According to the orthodoxy, intentional mental states are individuated by their propositional contents. These contents, in turn, have correctness conditions that specify what must be the case for the intentional mental state in question to be correct. Thus, at least part of what is required in providing an account of mental contents is to identify the facts that constitute their conditions of correction (Sultanescu 2015). Also, it is usually accepted that mental contents are aspectual or have modes of presentation: i.e., they represent certain entities under some specific aspects (Searle 2004; Crane 2001). Consequently, specifying the correctness conditions of mental contents will involve not one but two different tasks: i) identifying the item or entity that this mental content is about (the reference), and ii) determining the aspectual shape or mode of presentation of the item identified in i). Imagine, for example, that I think that *there is cup on the table*. A satisfactory account of the content of my thought must meet the following two constraints: it must explain what makes it the case that this thought is about the cup on the table, rather than about any other intermediate cause (i); and it must explain why my thought is about this cup *as a cup* or, in

other words, why its content involves the category *cup* instead of any of the other categories that can be adequately attributed to the thought's referent (such as *white object*, *ceramic object*, etc.) (ii) (Sultanescu 2015).

As mentioned above, neo-Pragmatists think that, in order to have thoughts with contents like, for example, *there is a cup on the table*, a creature has to be a participant in shared socio-cultural practices; practices based on social engagements and cultural devices, especially in linguistic tokens, which are the primary bearers of intentional content (Haugeland 1990; Hutto and Satne 2015). Thus, at least some of the constitutive facts that must take place so that even a relatively simple content like *there is a cup on the table* can be about the object cup *as a cup* are social facts.

According to Hutto and Satne (2015), neo-Pragmatist accounts face an “essential tension”: they are incapable of providing a non-circular account of mental content. Briefly, neo-Pragmatists claim that in order to have intentional mental contents, a creature needs to engage in social practices. Yet, being capable of participating in a socio-cultural normative practice requires intelligence and, more specifically, a host of cognitive abilities such as: being capable of learning, being capable of identifying and manipulating relevant elements that constitute the practice, having the ability to identify what is permitted or not in the practice, having the intention to take part in it, etc., (Satne 2016). Moreover, these cognitive abilities seem to presuppose the possession of mental contents. Thus, for example, it can be plausibly argued that in order to participate in a soccer game with others, I should be capable of identifying that the others intend *to play soccer*, I should have myself the intention *to play soccer*, I should be able to perceive *that these are the people who are playing*, that *this is the soccer ball*, etc. But, clearly, this involves entertaining at least some content-involving mental states. Neo-Pragmatists seem to fail, then, when it comes to providing a non-circular explanation of the origins and constitution of mental content.

Putting it in Hutto and Satne's terms:

“The puzzle is this: if all intentionality is of a piece and only derives from social practices, how is it possible that the sort of intelligent, recognitional capacities needed to explain participation in those social practices could be in place prior to their mastery?” (Hutto and Satne 2015, p. 529).

Now, Hutto and Satne propose a way to overcome this objection. Basically, their idea is that the problem of circularity dissipates if we posit not one, but two different varieties of intentionality. On the one hand, there is the full-blown content-involving intentionality of propositional attitudes, such as beliefs and desires. Following neo-Pragmatism, they claim that this kind of intentionality accrues only in socio-cultural normative contexts. On the other hand, there is a much more basic variety of intentionality, devoid of any content, which imposes only minimal requirements on its owners. In order to have this second variety of intentionality, a creature only needs to be capable of selectively *responding* to the world through targeted responses.¹ These targeted responses, in turn, must be

¹ By introducing this dichotomy, Hutto and Satne clearly depart from the orthodox assumption that intentionality essentially involves “having semantic content.” See also Hutto and Myin (2013) and Hutto and Myin (2017) for an extended defense of such view, according to which “mentality is not at root content-involving” (p. 11).

grounded in the history of past organisms of the creature's species. Since such responses are *directed* towards certain entities, it seems justified to claim that they exhibit some sort of intentionality. But, at the same time, it is important to keep in mind that these are non-representational responses to an object *simpliciter* (Rowlands 2015). In other words: these are responses that do not depend on the representation of the object in question under some specific aspect or other.

Hutto and Satne claim that the addition of this new variety of primitive intentionality will help explain “how it is possible that content could arise in the natural world” (p. 530). They also seem to think that this notion will play a crucial role in providing an account of the “natural origins of content” which delineates “a story of the transition from a world devoid of intentional states to a world abounding on them” by “invoking primitive forms of intentionality meant to play the role of intermediate stages” (Sultanescu 2015, p. 642). Thus, by positing a primitive intentional capacity that does not presuppose the possession of mental contents, it will be possible to account for the transition from a world lacking intentionality to a world in which social creatures immersed in shared social practices display full-blown intentionality. Their idea seems to be the following: in virtue of having primitive intentionality, young infants are able to enter into our socio-cultural world, gradually getting involved in our normative practices and, as a result of it, developing a capacity to represent other things and to entertain full-blown intentional mental states. If such an account is successful, the threat of circularity that has haunted neo-Pragmatism will disappear.

I will not discuss here whether this notion of primitive intentionality is all we need to explain in a non-circular way how infants lacking mental contents can become engaged in shared normative social practices and how, as a consequence, they come to have intentional contents (though see Rowlands 2015; Sultanescu 2015 for objections to this proposal). Rather, I would like to focus on the other difficulty, previously mentioned, that also imposes a significant problem for neo-Pragmatist accounts: the objection from animal minds.

As Haugeland (1990) himself admits, neo-Pragmatist approaches confine intentionality exclusively to social (and linguistic) beings, leaving no room for the attribution of intentional mental states to non-human animals that do not take part in socio-normative practices. This is due to the fact that a creature lacking all understanding of social norms will be incapable of participating in social interactions and acting according to such norms. Now, empirical evidence from studies in animal cognition is difficult to reconcile with such a view. According to it, various species of non-human animals display a wide variety of cognitive and behavioral abilities whose explanation seems to require the attribution of intentional mental states –or at least some kind of mental representations– to these non-human species (Allen and Bekoff 1997; Clayton et al. 2006; Carruthers 2004; Dennett 1996; Saidel 2009; Bermúdez 2003; Hurley 2003; Ristau 1991; Dretske 2006; Okrent 2007). Such as they stand, then, neo-Pragmatist accounts of mental content would be incapable of integrating these lines of empirical research.

Now, as stated above, it seems that Hutto and Satne's notion of primitive intentionality purports not only to explain how infants come to have minds but also to provide a general account of “how organisms progressed from a primitive Ur-intentionality to content involving forms of intentionality” (p. 533). Thus, it can be argued that their notion should provide us with the theoretical resources needed to accommodate the empirical evidence on animal cognition, allowing us to place the kind of intentionality that their behaviors

reveal at some point in the transition that goes from non-intentional entities to full-blown human minds. In the following section, I will go back to their proposal in order to examine the problems that it faces precisely at this task.

2 Primitive Intentionality and the Objection from Animal Minds

In the last decades, an impressive amount of empirical evidence of complex and flexible behavior in many non-human species has been accumulated. Let us focus on two kinds of behaviors, amongst many, that have been considered particularly striking:

- a. *Tool use and construction behavior*: Invertebrates, birds, and a wide range of mammals show an impressive diversity of tool using behavior and even of tool manufacturing. Chimpanzees, for example, use leaves and twigs to soak up water (McGrew 2004), woodpecker finches use twigs or cactus spines to search for larvae in holes that they cannot reach with their peaks (Vauclair 1996; Lack 1947), etc. Moreover, the acquisition of these abilities seems to involve learning, and they can be applied with some flexibility in different contexts.
- b. *Food caching and recovery behavior*: Several bird species show a remarkable tendency to hide surplus food in different locations and to recover it later on, in times of need. According to many researchers, birds' impressive capacity to recover the hidden food relies on their ability to integrate what they remember about the food that they have previously cached. Some of them, like Clark's nutcrackers, harvest pine seeds during the late summer and hide them in thousands of subterranean caches. Later on, they recover the food, apparently, by remembering where it was previously cached (Balda and Kamil 1998) and the cache size (Möller et al. 2001). Other birds, like scrub jays, that harvest both perishable and non-perishable food, face an additional task: they have to learn which kinds of food degrade and they need to update this information based on ecological conditions, on how long ago they have cached each particular item, etc. (de Kort et al. 2006).

The kind of evidence mentioned in a) and b) shows that some animals are capable of giving flexible, yet specific and adaptive, responses to different objects and properties in their environments. The question that we should ask ourselves now is: can we explain this kind of behaviors by adopting Hutto and Satne's proposal? Since these animals do not have a language and, presumably, neither participate in social practices nor follow social rules (at least in a relatively robust way), Hutto and Satne would be forced to entirely explain their behavior in terms of their content-free intentional relations with different objects and events in the world.² But, can we adequately explain the kind of behavior

² There is nowadays an interesting debate on whether some particularly clever non-human species – like apes— are sensible to some kinds of social norms. Defenders of the claim that apes have some basic, and implicit, understanding of social norms include Andrews (2009, 2013), Sultanescu and Andrews (2013) and von Rohr et al. (2011). Amongst skeptics, one may mention Slingloff and Moore (2018) and Rakoczy (2015). Now, even if we accept that some species do have social norms in a relatively robust sense, the problem of animal minds keeps its relevance, since there are still many other species that, for all we know, are not sensible to social norms and, still, show complex behavioral patterns that *prima facie* seem to call for explanations in terms of intentional mental contents.

mentioned in a) and b) merely by attributing to the agents a content-free variety of primitive intentionality? Could the capacity to give targeted responses to certain objects, events, properties, etc., suffice to explain these responses?

One could try to offer an affirmative answer to these questions on the following grounds: the kinds of behaviors mentioned in a) and b) may all be made possible by “special-purpose” non-representational mechanisms, triggered by specific features in the environment. They may be like the fly-snapping responses of frogs to any small, black, moving object: a rigid content-free, targeted behavior.

But, is this really so? Let me take a first stab at arguing that this is not a viable interpretation of the kind of empirical evidence presented in a) and b). As mentioned in section 1, it is a widely accepted claim that mental contents have both a reference and an aspectual form, or a mode of presentation. Hutto and Satne’s primitive intentionality, on the contrary, is a non-representational capacity to give specific targeted responses to certain objects *simpliciter* (i.e., a response to these objects understood as not falling under any particular mode of presentation) (Rowlands 2015). Thus, it involves a relationship towards certain objects that does not depend on any specific way of thinking, characterizing, or representing them.

If this is so, when an organism A gives a targeted response to an object O, it can be claimed that O is the referent of A. But, given that this is a response to O *simpliciter*, it should not depend on, neither be affected by, A’s previous representation of any specific properties, or features of A. On the contrary, it is to be expected that, whenever A registers O, or detects O’s presence, A should give exactly the same targeted response towards O, no matter which are the properties of O that are salient or manifest at that moment, which are the previous knowledge and representations regarding O that A has acquired, etc.

Of course, there will be at least some “excusing conditions” to this general claim that should be considered.³ A may somehow be prevented from giving her expected response to O, because of some physiological failure. This will be the case, for example, if A’s typical response in front of O consists in running, but, at the present moment, A has broken one of her legs. Another possibility is that somehow the context changes in a way that precludes A from giving its usual response to O. This is what happens with the creature that normally responds to food by grabbing it, but today ignores it because she is too busy running away from a predator. In cases like this, A’s behavior should not count as a reason against claiming that, *ceteris paribus*, A always gives a specific targeted response to O *simpliciter*. Clearly, it will be necessary to provide a principled way of distinguishing between admissible and non-admissible excusers. I will not undertake such a task here. Nevertheless, and very roughly, I would like to suggest that all admissible excusers should at least have something in common: they must all be situations in which certain non-representational variables somehow impede or trump the usual targeted response from taking place.

Let us imagine, on the contrary, that A displays the capacity to give different responses to the same object O, in different circumstances, even when it is not possible to identify relevant excusing conditions to explain that deviance. Thus, for example, imagine that A

³ The notion of “excusing conditions” was originally introduced by Eric Schwitzgebel (2010) in the context of discussing his dispositional account of beliefs. Nevertheless, it seems to me that it can be safely extended to the present debate.

responds to O by eating when hungry, but also uses it as a weapon, throwing it against her enemies when she is in the middle of a fight. In cases like this, it can be conjectured that A is representing O in different ways, in different situations, and that this is what causes the divergence in her responses.

If the previous line of thinking is correct, the explanations that appeal to primitive intentionality, understood as the capacity to give targeted responses to objects *simpliciter*, will be clearly valid in the case of those animals that behave as “mere reactors” (Duhau 2010). These are creatures that (leaving aside excusing conditions) always give the same specific response R_1 to the same object O. Moreover, they are incapable of learning anything new about O and, consequently, their responses to O remain rigid and inflexible through time. Since their behavior towards O remains the same (at least in normal conditions), it is unnecessary to postulate neither that they are representing O under any specific mode of presentation, nor that they are harboring a propositional content, amongst other possible ones, in which certain property is predicated of O. Such invariance of behavioral responses strongly suggests, rather, that they are merely responding to this object *simpliciter*.

Yet, there is also another possibility to contemplate. Hutto and Satne’s account may also be capable of successfully dealing with creatures that in T_1 rigidly respond to O in a way R_1 but, later on, as a result of some non-representational process, change their behavior and rigidly respond to O in way R_2 in T_2 . The process that leads to the behavioral change must be non-representational, because Hutto and Satne assume that the creatures in question are only endowed with a primitive variety of content-less intentionality. But, if there is a plausible story to tell along these lines, they can claim that those creatures that change one kind of rigid response to an object O for another have simply replaced their original disposition to give a specific response to O for a different one. This could be, for example, the case of those animals that initially respond to a piece of food by eating it, but once they have gotten sick several times after consuming the food in question, start avoiding it in a systematic way.⁴ Anyway, it should be noted that, in these cases, the animals are still incapable of giving more than one rigid targeted response *per* object in a given period of time. They may eventually change their original response for a different one, but they never acquire the versatility to respond to the same object in one way or the other.

However, even if the notion of content-free primitive intentionality is well-fitted to describe the behaviors mentioned above, empirical evidence like the one presented in a) and b) suggests that many non-human animals are capable of much more than rigidly reacting to certain objects in only one specific way in a given period of time. Let us consider the case of tool use by woodpecker finches. In their natural environments, these birds frequently use twigs or cactus pines to search for larvae in holes that they cannot reach with their peaks. This is not a rigid reaction to the presence of cactus pines or twigs, since when these objects do not have the appropriate size, they modify these tools by shortening the twigs, or by breaking the side twigs that would prevent insertions into holes (Tebich and Bshary 2004). Now, this suggests that birds are not merely giving targeted responses to certain objects *simpliciter* – twigs or pines— but rather responding to these

⁴ In this case, the process that leads to the relevant behavioral change is a process of associative learning, and the example will work only if we assume that the learning process involved is non-representational. I will make that assumption here, since it only makes Hutto and Satne’s account more difficult to challenge.

objects in different ways when their properties differ: using them directly as tools when their size is appropriate and previously modifying them when they are not. Moreover, the specificity of their flexible behavior suggests that they are also representing these twigs as having certain properties (a certain size or shape) that make them appropriate tools for a given task (or not). Similarly, nutcrackers show a wide spectrum of complex and flexible behaviors towards pine seeds. While harvesting them, for example, they distinguish between “edible” and “inedible” seeds by “bill clicking” and “bill weighing” (Balda et al. 1996). Nutcrackers cache their seeds on their breeding territories and, also, on “communal” caching areas where many birds intermix caches. Now, when caching their food in these communal areas, these birds act very secretively, often perching in a tree for some minutes and peering down intently before diving to the ground. Later on, they recover the cached food by using spatial memory based on visual landmarks (Balda and Kamil 1998). Thus, nutcrackers seem not only capable of acting in specific ways toward seeds (harvesting, caching and recovering them), but also, according to the versatility of their responses, of detecting and representing specific properties about them (whether they are edible or not, whether they were hidden on communal ground, whether they were hidden next to a specific landmark, etc.).

I think that there is still a way in which defenders of primitive content-free intentionality may respond to this objection. So far, I have suggested that, in the empirical examples examined in a) and b), animals were giving different flexible responses to *the same objects*, because they were representing those objects *as having different properties*. And, clearly, in that case they would have to be capable of entertaining mental states with intentional contents. However, there is an alternative explanation of the evidence: these animals may be just giving specific targeted responses to different *conjunctions of properties*. Even though Hutto and Satne do not discuss this option explicitly, it seems to be compatible with their proposal, since targeted responses to properties are not guided by any representation of these properties. Moreover, positing a capacity to respond to conjunctions of properties, rather than to objects, might give them more powerful explanatory tools to deal with the empirical cases that we are discussing.

Let us see, for example, how the notion of a targeted response to properties could be used to explain the behavior of woodpecker finches. Basically, it could be claimed that when a specific bundle of properties is instantiated in a particular twig – i.e., when the twig has the kind of length L_1 and size S_1 (which are, precisely, the length and size that turn it into an appropriate tool to search for larvae)—, the birds respond to those properties by using the twig to search for larvae. But, when the bundle of properties is a different one, they act on the twig, modifying it until it has the appropriate set of features (the ones that are required to search for larvae). One could offer, additionally, a similar account of the capacity of Clark’s nutcrackers to distinguish between edible and non-edible seeds, or between seeds that are hidden in a communal area and seeds that are not. In each case, it might be argued, these birds behave differently because they are responding to a (slightly) different bunch of properties in the world.

It might be objected that this kind of explanation seems rather arbitrary and ad hoc, even for rather simple cases, like the one of the woodpecker finches described above. Very likely, many manifest properties will change in different encounters of woodpecker finches with those twigs that they will end up using as tools: the shape of the twig, the number, color and shape of its leaves, their spatial dispositions, etc. Thus, it seems legitimate to ask: do we have good reasons to believe that there is a common core of

properties that are present in all the situations in which an animal gives the same kind of response, so that we can explain the latter as due to the former? Furthermore, how can we establish which are the properties that the woodpeckers are systematically responding to in all these situations, and which ones are they ignoring?

When we conjecture that an animal is capable of thinking about *objects as having properties*, we acquire a plausible and elegant way to answer these questions. When woodpeckers modify a twig by cutting its branches, or by shortening it, we can claim that a representation of *this twig as being too large*, or something like it, is guiding their behavior. Thus, as long as the birds have categorized the object out there as a twig, and as having the appropriate size and length, one may expect the adequate behavior to follow, no matter which other properties may have changed from one situation to another.

Even if we admit that the birds in question do not share our exact concepts of *twigs* or of *being too large*, we can hypothesize that they have some categories that refer to twigs and to their size and that perform analogous functions, at least in this kind of practical situations, to the ones played by our concepts. And, even if we are not sure about the exact concepts and categories that these animals are using, by assuming that they think about these particulars and their properties using categories that are similar to ours, we already gain some capacity to make accurate predictions about their future behavior. We can anticipate, for example, that they will respond to a twig that they have never seen before in the same way in which they have responded to other twigs, or we can predict that if we add false plastic leaves to a twig, so that if it ends up being too large, they will peel them away, etc.

In brief, attributing semantic contents to these birds gives us some initial explanatory and predictive leverage. Furthermore, it allows us to understand why they give the same responses whenever they find a twig, even if other properties of different objects near them change (because they are responding to twigs and their properties and not to these other objects), why they ignore irrelevant differences between twigs (because they are only interested in the size of the twigs and not, for example, in its color or texture), etc. But, if we claim, instead, that the woodpecker finches are merely responding to “different bundles of properties”, without further specification, we are left without any theoretical tools to answer these kinds of questions.

Finally, if we are not convinced by these preliminary considerations, it is still possible to develop a second, and stronger, line of argument to show the need of attributing to some non-human animals the capacity to represent particulars as having certain properties instead of merely crediting them with targeted responses to conjunctions of properties. This new argument focuses on a different kind of empirical evidence: evidence that (some) non-human animals are capable of learning and integrating information, acquired in the past, about the properties of particular objects, and of using that information when those properties are not salient or manifest anymore.

Let us consider the following example: in a well-known study, Clayton et al. allowed a group of scrub jays to cache worms and peanuts in visuo-spatially distinct ice-cube trays. Each tray consisted of two parallel rows of ice-cube moulds filled with sand. In the experiment, the birds cached one food – crickets – in one side of a tray, and another food—peanuts— on the other side of the same tray. On some trials, they were allowed to recover both peanuts and fresh crickets after one day, whereas on other trials, the opportunity for recovery was delayed for four days, by which time the crickets had

decayed. Now, what happened was that the birds rapidly learned to search for crickets when fresh, and to search for peanuts, instead, when the time interval between caching and recovery was such that the crickets must have decayed. Thus, while on the first two trials the majority of the birds searched first for crickets, even after the 4-day retention interval, by the third trial, all birds switched their preference and were searching in the peanut side first after the long-interval. Despite this switch, the majority of the birds continued to direct their first search to the cricket side after the 1-day retention intervals.

Now, it seems that jays have acquired different kinds of information during their training experience. First, they have learnt about the temporal decay of crickets (and not of peanuts). More specifically, they have learnt that crickets are fresh at retention intervals of less than 4 days, but not after longer intervals. But, in order to behave as they do, they must also have a memory representation of: i) *what* they cached on the trays (peanuts or crickets), ii) *where* they cached each kind of food, and iii) *when* they did this. In fact, researchers have dubbed these as “what-where-when” representations (Cf. Clayton and Dickinson 1998; Clayton et al. 2006).

To my mind, these results also show that scrub jays are integrating information about different properties of the food they have hidden. Moreover, much of this information is about properties that they have experienced in the past, but which are not actually manifested. Thus, for example, in order to look for the crickets after the 1-day retention interval (but not after the 4-day retention interval), they must remember *that some crickets where cached on the left side of the tray, that they were left there 1 day ago, and that crickets are fresh after a 1-day interval*. But, there are two interesting things to notice here. First, in order to successfully deal with the experimental task, the jays must not only represent the place in which the items of food are hidden, the period of time that has passed since caching them, and their edibility or inedibility. They must also represent all these properties as properties of the *same particular crickets* or of the *same particular peanuts*. This is what allows them to integrate all this information together and leads them to act as they do towards the crickets and peanuts. Secondly, the jays experienced in the past that crickets and peanuts had different properties. But, since the latter are not presently manifest, the birds cannot perceive where the items of food are hidden at present, or whether they are still fresh or not. They must remember this information.

In brief, it seems that, in order to behave as they do, scrub jays need to have some capacity to acquire, store, and apply information about some properties of crickets and peanuts. In other words, they must not only be sensible to these properties when they are manifest in their surroundings, but they must also be capable of forming and carrying stable representations of them that can be used later on. In addition, it can be argued that they must also represent all these different properties *as properties of the same kind of particulars* – peanuts or crickets—. They must have some understanding of the fact that it is *crickets* that were hidden, for example, on the left side of the tray, that were cached one day ago, that are still fresh, etc.

Merely attributing a capacity to give targeted responses to conjunctions of properties does not seem to allow us, on the contrary, to explain the behavior of scrub jays during the experiments. In particular, one will not find it easy to explain the change in the scrub jays’ behavior once they reiteratively experience that crickets decay after the 4-day interval. As far as I can see, one may at most sketch the beginnings of such an explanation by claiming that scrub jays have learnt to respond to the manifest properties

of a certain patch of land (where they originally hid the crickets) by preferentially searching for food there. That would allow explaining how the birds, which normally prefer crickets to peanuts, originally searched first on those places where they had hidden the former. But, the problem is that experiments show us that, after experiencing that crickets are rotten after a 4-day interval, while they remain fresh after a one-day interval, scrub jays change their original behavior and they search for peanuts first after every four-day interval. How do we explain this change by positing only a tendency to give targeted responses to bunches of properties? Presumably, the manifest properties that the birds can be related to when they are searching for the hidden food have not changed in any relevant way. All that has changed are their past experiences of crickets and peanuts. It seems then that an account in terms of mere targeted responses to properties will get stuck at this point. Contrarily, attributing to these birds the capacity to represent some properties of crickets and peanuts will allow us to provide a swift explanation of their behavior.

3 Primitive Intentionality and REC: Further Developments

In the previous sections, I have focused on Hutto and Satne's notion of primitive intentionality as they develop it in their paper "Natural Origins of Contents" (2015). However, the notion has been further defended and enriched in a later book by Hutto and Myin (2017): *Evolving Enactivism: Basic Minds Meet Content*. One may legitimately wonder, then, whether it wouldn't be possible to find in this more recent work additional resources to successfully deal with the objection from animal minds.

Hutto and Myin's book is an explicit attempt to defend a Radical Enactive and Embodied account of Cognition (REC). According to REC, cognition should be understood as a kind of situated, enactive, and embodied activity. In addition, REC advocates hold that basic forms of cognition lack contents and consist only in "organisms actively engaging with selective aspects of their environment in informationally sensitive, spatiotemporally extended ways" (Hutto & Myin 2017, p. ivx).

Strictly speaking, Hutto and Myin do not attempt to develop *one* positive account of cognition in their book. Rather, they turn to a range of extant theories of cognition – including the Predictive Processing Account of Cognition, Ecological Psychology, Sensorio-motor Theory and Autopoietic Adaptive Enactivism— to see whether, with some rectifications, each of them can be brought under REC's umbrella. Clearly, even a rough examination of these different proposals and their potential usefulness to explain challenging evidence on animal cognition exceed the limits of this paper. These "rectified" theories could eventually provide ways of overcoming the problem of animal minds that I might overlook here. Nevertheless, in what follows, I would like to briefly and tentatively review some ideas present in Hutto and Myin's book that seem especially relevant to the problem in question, and to give some reasons to remain skeptical about their prospects to explain the kind of evidence on animal cognition discussed in section 2, without positing any kinds of intentional contents.

One central point to focus on is Hutto and Myin's detailed attempt to provide an account of their new variety of primitive intentionality in terms of a "rectified" version of Millikan's teleosemantic theory of content. Their basic idea is that one can understand the specific kind of world-directness exhibited by those creatures endowed only with

primitive intentionality by developing an anti-representationalist variety of teleosemantics—which they call “teleosemiotics”—and putting it to a new theoretical use: explaining and understanding basic cognition.⁵

One key advantage of teleosemiotics, they claim, is that it provides us with the resources to explain primitive intentionality as a relation between a creature and a genuine object, instead of a relationship between a creature and a mere thing. As seen above, Hutto and Satne think that primitive intentionality consists in the capacity to give targeted responses to objects. To focus on one famous example, frogs lashing out their tongues are giving a targeted response to flies. But the problem is that any creature may *target* one object but, eventually, end up responding to other different things in their environments. Frogs, for example, may occasionally lash their tongues to other stimuli that are similar enough to flies. But, then one may ask: how can it be established which the object that is being targeted by a creature is as opposed to the many things that she may actually respond to?

Teleosemiotics provides us with the following answer: we may claim that a creature—for instance, a frog—is giving a targeted response to a specific kind of object—like flies—because these creature’s “perception-action routines were forged through a long process of selection by consequences in order to get flies, and not anything else into its belly” (Hutto and Myin 2017, p. 115). Moreover, it allows us to explain *why* the targeted response of the frog is a response to *flies*: basically, this is so because eating flies—and not any other object—is what helped the frog’s ancestors to meet their survival needs.

So far so good. Teleofunctionalism may provide us with an adequate way of delimiting which the objects that a creature’s responses are targeting are, even if she is not representing those objects in any specific way. And positing this kind of targeted responses to objects *simpliciter* may be all we need to account for some varieties of basic cognition and behavioral patterns. Nevertheless, this move falls short from providing what is needed to overcome the problem of animal minds. As I have previously argued, based on empirical examples, what needs to be explained is not only the capacity of animals to respond to specific objects, but, rather, how their responses vary flexibly and adequately, not only in virtue of the objects that they are interacting to, but also in virtue of these objects’ properties (especially given that these properties may vary from one object to the other, may be manifest in some circumstances but not in others, and may change in time, even in the same object).

Admittedly, Hutto and Myin do take some steps that might help advocates of primitive intentionality to explain some of these features of animal behavior. They begin by claiming, turning once again to teleosemiotics, that despite being forged by evolution, the ways of responding to the world of some animals need not be rigidly fixed and hardwired. On the contrary, Mother Nature may have also provided them with abilities to learn new ways of responding to the environment through their individual experiences (cf. Hutto & Myin, p. 117). Moreover, following Bruineberg and Rietveld (2014), they tell us that it is also possible to provide a content-free explanation of how organisms can be

⁵ The original aim of teleosemantics, Hutto and Myin claim, was to offer a naturalistic biological account of mental content. In a nutshell, teleosemantic theories understand mental representations as inner states with the biological function of enabling organisms to keep track of items in their environment. An inner state has the function of representing an item if and only if it was selected for that task (i.e., if performing that function helped these animals’ ancestors to reproduce or survive). Hutto and Satne (2015), Hutto and Myin (2013), and Hutto and Myin (2017) argue that teleosemantics is doomed to fail at this ambitious original task.

capable of adjusting and adapting their behaviors in a manner that allows them to respond in flexible ways to the multiple possibilities of action that their environments afford in different circumstances. Provided that such accounts succeed, at least to a certain extent, it might be possible use them to explain some of the flexible and dynamic responses of animals to their changing environments.

But this is not all. Hutto and Myin also argue that it is possible to understand basic kinds of imagining and remembering as consisting only of content-free embodied activities and reenactments. Interestingly, the preliminary content-free accounts of memory and imagining advanced by them do make use of some theoretical posits like mental “stands-in”, “surrogates” and “simulations” that many philosophers and cognitive scientists have considered to be of a representational nature. However, Hutto and Myin argue that merely positing cognitive states or processes that play the role of stands-in, or surrogates, of items in the external environment – even if engaging with those surrogates involves exploiting some correspondences between them and the items modeled— does not suffice to conclude that these states or processes have representational contents.⁶ In order to be regarded as contentful, they must also have truth or correctness conditions that play a genuine role in the explanation of the creature’s behavior.

Admitting that a creature that only has primitive intentionality may engage in some sort of simulation processes, or make use of surrogates, in order to guide her behavior, might be of help when it comes to explaining a challenging feature present in the behavior of many non-human animals: its detachment or independence from present stimuli. In effect, we have evidence of different animals’ capacity to respond not only to objects or properties that are present in their environments, but also to respond in ways that seem to be sensitive to things that happened in the past, to things that have not yet happened, etc. Orthodoxy would recommend positing content-involving processes of memory, imagination, simulation, etc., to account for such responses. But, if Hutto and Myin claims are correct, and it is possible to develop sound content-free explanations of how animals remember, imagine, etc., then mental contents would seem explanatory gratuitous to deal with such evidence.

However, I think that there are still several reasons to be cautious regarding the potential of Hutto and Myin’s theoretical developments to help advocates of primitive intentionality to overcome the problem of animal minds. In the first place, what Hutto and Myin offer are still incipient and general proposals on different aspects of basic cognition that they do not extend to deal with the kind of empirical evidence on animal cognition discussed above. There is, then, still much work to be done if one wants to defend, based on their ideas, that the cognition of non-human animals is entirely content-free.

Secondly, at least as I understand them, Hutto and Myin’s recent attempts to develop content-free accounts of different cognitive phenomena may provide us with helpful tools to deal with (at least some varieties of) the flexibility and detachment from immediate stimuli that the behavior of many non-human animals display. This is not trivial, since several philosophers have focused on these properties – amongst others— in order to justify their attribution of mental contents to non-human animals.⁷ However,

⁶ It is important to keep in mind that according to REC defenders, even if imaginers engage with surrogate mental models, such an activity will be fundamentally embodied and enactive. Consequently, imagining will be strongly constrained by the actual ways in which the organisms engage with the objects modelled (and vice versa). Cf. Hutto and Myin (2017, p. 200).

⁷ See, for example, Saidel (2009), Camp (2009), Bermúdez (2003) and Bennett (1990).

it seems to me that this is not sufficient to overcome the problem of animal minds. Let me anticipate my argument.

There is a wide consensus amongst philosophers that mental contents have several constitutive features: they have truth or correctness conditions, they may represent false or non-existing things, and they usually represent these things as being one way or another. Now, I will try to show that the behavioral patterns of some animals show not only flexibility and distance from immediate stimuli, but also two other characteristics that I will call *sensitivity to error* and *sensitivity to the aspects of objects*. As I will argue, the explanation of those further traits seems to require positing states or processes with two of the core features of contents: correctness conditions and aspectuality. Consequently, a complete explanation of extant evidence on animal behavior will still involve positing mental contents.⁸

Let us begin by focusing on sensitivity to error and correctness conditions. In a recent paper, that focuses on when non-human animals can be literally credited with intentional mental contents, Knoll and Rey (2017) begin by admitting that ascribing mental contents to animals is part of the best explanation of their behavior when "properties that are arguably constitutive of being and intentional representation... are essential to the explanation" (p. 14). It is widely agreed, they go on arguing, that having correctness conditions is one of the constitutive features of intentionality. Then, one of the requirements to be justified in attributing mental contents to non-human animals is that correctness conditions play a genuine role in the explanation of their behavior.

Imagine, now, that some animals are insensitive to their errors and, consequently, lack all capacity to recover from them. Even assuming that these animals had cognitive "stand-ins" that were correlated with objects in their environments and that they guided their behavior based on those correlations, if these correlations broke and led them to act inadequately, they would be incapable of correcting their cognitive states and their actions. It seems that, in such a case, the correctness conditions of those putative stands-in would play no role in the explanation of their behavior.⁹ But if, on the contrary, our animals could somehow recover from their errors, mental contents would become relevant in the explanation of their behaviors. Why? To my mind, the main reason is the following one: because what explains their capacity for recovery is their capacity *to respond* to the fact that they are in cognitive states whose correctness conditions do not obtain, and to *correct* those states (so that they guide these animals' behavioral responses in a more adequate way). Or, turning now to a representationalist' way of talking: they are sensitive to their representational errors, and they are able to correct them (by updating or modifying their representational contents).¹⁰

⁸ Hutto and Myin also stress that an attribution of mental contents to an organism is legitimate when core features of contents – like correctness conditions— have a genuine role to play in explaining its behavior. I think this demand is fair, but I will argue that it can be met.

⁹ Knoll and Rey (2017, p. 16-17) claim that a good example of the incapacity to recover from errors is the ants' impossibility of correcting their paths back to their nests after having been displaced.

¹⁰ It is important to bear in mind, however, that in order to be sensitive to their representational errors, animals need to respond appropriately to those errors (by correcting them). But, in order to do so, they need not represent, or conceptualize, neither their representational errors nor the correctness conditions of their mental states as such (Glock 2000; Knoll and Rey 2017; Danón 2011).

Focusing once again on empirical evidence, it is plausible to claim that there are non-human animals that enjoy this primitive capacity to recover from their errors and to update or adjust their representations about their environments. It seems plausible to argue, for example, that the scrub jays reared in captivity were erroneous when expecting to find the crickets that they had hidden four days ago to remain fresh. But, clearly, they were not insensitive to that error: when provided with relevant evidence through their individual experiences, they learnt better. Following our previous reasoning, we could argue that it is legitimate to attribute mental contents with correctness conditions to them in order to explain their change in behavior.

This first argument can be reinforced by showing that other core features of intentional contents also play a key role in the explanation of some non-human animals' behavior. Let me try, then, to establish that there is at least a second feature of intentional mental states that seems to be relevant when it comes to explaining some of the evidence provided in section 2: the *aspectuality* of content.

Different philosophers have argued, once and again, in favor of the necessity of positing aspectual mental contents in order to explain some behavioral patterns of different non-human animals (Gozzano 2007; Glock *forthcoming*; Knoll and Rey 2017; Danón 2016). The topic is complex and deserves a longer and more nuanced treatment than the one that it will possible to offer here. I would like, nevertheless, to give a rough argument in favor of such a claim. Much of what I am about to say recapitulates considerations already presented in the previous sections but, hopefully, it puts them under a slightly different light.

Out there in the world, objects have a myriad of different— an in many cases changing— properties. As has been discussed in the previous sections, it is possible to imagine different ways in which a creature may react to these objects and their properties. One possibility is for an animal to always respond to the presence of the same object in a way which is not affected by whether the animal in question focuses, registers, has epistemic access to, or takes into account some of these objects' properties instead of others. In such a case, it seems to be unnecessary to postulate that the animal's responses to the object are sensitive to some of the object's aspects (i.e., to some of its properties) instead of others. A second alternative is that an animal might be just responding in different stable ways to bunches of properties in the environment. Arguably, if such were the case, it would not be necessary either to credit these animals with a capacity to take some things, or objects, as being in certain ways, because they would not be capable of discriminating between a bunch of properties and the thing or object that has those properties.

Yet, there is also a third alternative: some animals might be capable of responding in a determinate way to an object, but only in virtue of the fact that this object has some specific properties. Thus, at least in principle, the presence of the object without the properties, or the presence of the same properties in a different object, would not trigger this specific response. In such a case, these animals would not be neither responding to objects, nor to bunches of properties. They would be responding to some objects in a specific way, but only because they are capable of *taking these objects to be in a certain way*. But then, it seems that what enables animals to perform these specific patterns are some aspectual mental contents in which things are represented as being in certain ways or as having specific properties.

Once again, it seems to me that we have good evidence of animals who are capable of satisfying such a requirement. The scrub jays studied by Clayton and colleagues are probably a good example of it. Clearly, they can learn to respond to the same pieces of food in a way that depends on their properties. That is why they preferentially search for crickets when those insects have only been underground for a brief period of time and remain fresh, but they change their preferences and search first for peanuts when the crickets have been hidden for a longer period of time and have decayed. But, when doing so, they do not seem to be merely reacting to some properties out there in their environment. One reason to think that this is not the case is that the birds in question have no perceptual access to the relevant properties (e.g. the actual states of edibility of the hidden pieces of food). And, more importantly, because in order to somehow track whether what lies underground is fresh or rotten they need to remember what they have hidden (since crickets decay, but peanuts do not) and where they have hidden it. It seems more plausible to claim, then, that they are responding to some specific objects – the particular pieces of food that they have hidden in particular moments of time and in specific places— but also in a way that takes into account these objects’ varying properties. In brief: they seem to be taking these crickets or peanuts as being fresh or rotten. But, if this is correct, aspectual mental contents do seem to play a genuine role in the explanation of this kind of complex behavioral patterns.

Recapitulating: advocates of primitive intentionality may take some steps further towards explaining different features of non-human animals’ cognition and behaviours by leaning on recent developments by Radical Enactivists, like Hutto and Myin (2017). However, as far as I can see, this notion will still fall short from explaining many interesting features of animals’ behaviour, like their sensitivity both to their errors, and to the different aspects of particular objects. *Pace* Hutto and Satne, intentional mental contents still seem to play a relevant role when it comes to explaining what animals can do.

4 Conclusions

In this paper, I have argued that neo-pragmatist accounts of mental content face what I have called the objection from animal minds. Very roughly, due to the emphasis that they put in the social character of mental content, neo-Pragmatists end up exclusively confining intentionality to social (and linguistic) human beings, and rejecting the attribution of intentional mental states to those non-human animals that lack language, and do not take part in socio-normative practices (at least in a strong sense of the term). But, this is in clear tension with the fact that, as argued above, the best way of explaining an impressive amount of empirical evidence on animal cognition seems to require attributing intentional mental states to different non-human species (including some solitary species that seem capable of displaying impressively flexible and innovative behavioral patterns).¹¹

¹¹ See Holekamp (2007), p. 67.

Later on, I examined the explanatory potential of the notion of primitive intentionality to help neo-Pragmatists deal with the problem of animal minds. I claimed that Hutto and Satne's new variety of intentionality could be understood in two ways: i) as a capacity to give targeted responses to certain objects in the creature's environment, or ii) as a capacity to give targeted responses to sets of properties present in the creatures' surroundings. Thus, if there is empirical evidence of non-human animals' behavioral responses that cannot be explained either by i) or ii), it is possible to conclude that the notion of primitive intentionality will not suffice to overcome the problem of animal minds.

In section 2, I examined some evidence that, as far as I can see, cannot be adequately explained by i), and some other evidence that neither can be explained by i) nor by ii). On the contrary, I have suggested that the most plausible account of this kind of evidence requires attributing to these animals mental contents that represent particular entities as having certain properties. In section 3, after examining some recent developments in REC, I gave further reasons to keep endorsing this claim.

Does this mean that we must attribute them the same kind of contents that we attribute to linguistic and socially competent human animals? Not necessarily, since there are theoretical alternatives that – despite their differences— converge in attributing to non-human animals the basic capacity to represent some particulars and to characterize them as having certain properties, while, at the same time, they refrain from attributing to them contents with the same kind of sentential structure, logical form, and full-fledged recombinability that we usually ascribe to the contents of the mental states of adult human animals (see, amongst others, Sellars 1981; Millikan 2006; Burge 2010a, b). If we follow this line of thinking, a promising alternative emerges. We can enrich our repertoire of kinds of intentionality by adding to the full-blown intentionality of propositional attitudes not only the contentless intentionality posited by Hutto and Satne, but also another variety of primitive intentionality that can be shared with non-human animals and pre-linguistic infants and that is cognitively less demanding than full-blown intentionality, but still involves semantic contents *that represent some things as being in certain ways*.

Although I will not be able to develop such a proposal to any extent here, I would like at least to point out some of its expected virtues. Firstly, such a theoretical approach may help us develop an account of animal intentionality that respects the analogies that we may find between other species' behaviors and ours while, at the same time, preserving the idea that there are important differences between our mental life and theirs. Secondly, it should allow us to overcome the problem of animal minds, while retaining the basic neo-Pragmatist idea that engaging in social practices is, indeed, necessary to have (some kind of) intentional mental states. Nevertheless, there is a cost that neo-Pragmatists would have to pay in order to follow this route. They would have to limit the scope of their original claims about the dependency of mental contents on social and linguistic practices to those mental contents that are exclusive of human animals. Or, in other words, they would have to admit that even if it is true that linguistic human beings have a unique kind of intentional mental contents which only emerges in linguistic, social, and cultural contexts, there are other kinds content-involving cognitive states and processes that extend far beyond our species.

References

- Allen, C., & Bekoff, M. (1997). *Species of mind: The philosophy and biology of cognitive ethology*. Cambridge: MIT Press.
- Andrews, K. (2009). Understanding norms without a theory of mind. *Inquiry*, 52(5), 433–448.
- Andrews, K. (2013). Ape autonomy? Social norms and moral agency in other species. In K. Petrus & M. Wild (Eds.), *Philosophical perspectives on animals: Minds, ethics, morals* (pp. 173–191). Bielefeld: Transcript Verlag.
- Balda, R. P., & Kamil, A. C. (1998). The ecology and evolution of spatial memory in corvids of the southwestern USA: The perplexing pinyon jay. In R. P. Balda, I. M. Pepperberg, & A. C. Kamil (Eds.), *Animal cognition in nature: The convergence of psychology and biology in laboratory and field* (pp. 29–64). San Diego: Academic Press.
- Balda, R. P., Kamil, A. C., & Bednekoff, P. A. (1996). Predicting cognitive capacity from natural history: Examples from four species of corvids. In V. Nolan & E. D. Ketterson (Eds.), *Current ornithology* (Vol. 13, pp. 33–66). New York: Plenum Press.
- Bennett, J. (1990). *Linguistic behaviour*. Indiana: Hackett Publishing Company.
- Bermúdez, J. L. (2003). *Thinking without words*. Oxford: Oxford University Press.
- Brandom, R. (1994). *Making it explicit: Reasoning, representing and discursive commitment*. Cambridge: Harvard University Press.
- Brandom, R. (2000). *Articulating reasons: An introduction to inferentialism*. Cambridge: Harvard University Press.
- Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8, 599.
- Burge, T. (2010a). *Origins of objectivity*. New York: Oxford University Press.
- Burge, T. (2010b). Steps towards origins of propositional thoughts. *Disputatio*, 4(29), 39–67.
- Camp, E. (2009). Putting thoughts to work: Concepts, stimulus-independence and the generality constraint. *Philosophy and Phenomenological Research*, LXXVIII, 275–311.
- Carruthers, P. (2004). On being simple-minded. *American Philosophical Quarterly*, 41(3), 205–220.
- Cash, M. (2008). The normativity problem: Evolution and naturalized semantics. *The Journal of Mind and Behavior*, 29(1/2), 99–137.
- Clayton, N., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, 395, 272–274.
- Clayton, N., Dickinson, A., & Emery, N. (2006). The rationality of animal memory. Complex caching strategies of western scrub jays. In S. Hurley & N. Mathews (Eds.), *Rational animals?* (pp. 197–216). Oxford: Oxford University Press.
- Crane, T. (2001). *Elements of mind*. Oxford: Oxford University Press.
- Danón, L. (2011). Normatividad doxástica en animales no humanos. *Analítica*, 5(5), 9–28
- Danón, L. (2016). Opacidad referencial y atribución intencional a animales sin lenguaje. *Principia*, 20(2), 143–164.
- de Kort, S. R., Tebich, S., Dally, J. M., Emery, N. J., & Clayton, N. S. (2006). The comparative cognition of caching. In E. A. Wasserman & T. R. Zentall (Eds.), *Comparative cognition: Experimental explorations of animal intelligence* (pp. 602–618). New York: Oxford University Press.
- Dennett, D. (1996). *Kinds of minds*. London: Weidenfeld & Nicholson.
- Dretske, F. (2006). Minimal rationality. In S. Hurley & N. Mathews (Eds.), *Rational animals?* (pp. 107–116). Oxford: Oxford University Press.
- Duhau, L. (2010). Do bees really have concepts? A discussion of Carruthers' criteria for conceptuality. *Teorema*, 29(2), 125–134.
- Glock, H. J. (2000). Animals, thoughts and concepts. *Synthese*, 123(1), 35–54.
- Glock, H. J. (forthcoming). Animal minds: conceptual problems.
- Gozzano, S. (2007). The beliefs of mute animals. In M. Marraffa, M. D. Caro & F. Ferretti (Eds.), *Cartographies of the Mind. Studies in Brain and Mind*. Dordrecht: Springer.
- Haugeland, J. (1990). The intentionality all-stars. *Philosophical Perspectives*, 4, 383–427.
- Holekamp, K. J. (2007). Questioning the social intelligence hypothesis. *Trends in Cognitive Sciences*, 11(2), 65–69.
- Hurley, S. (2003). Animal action in the space of reasons. *Mind and Language*, 18(3), 231–256.
- Hutto, D., & Myin, E. (2013). *Radicalizing enactivism: Basic minds without contents*. Cambridge: MIT Press.
- Hutto, D., & Myin, E. (2017). *Evolving enactivism: Basic minds meet content*. Cambridge MA: MIT Press
- Hutto, D., & Satne, G. (2015). The natural origins of content. *Philosophia*, 43(3), 521–536.

- Knoll, A., & Rey, G. (2017). Arthropod intentionality? In K. Andrews & J. Beck (Eds.), *The Routledge handbook of philosophy of animal minds* (pp. 13–24). London: Routledge.
- Lack, D. (1947). *Darwin's finches*. New York: Cambridge University Press.
- McGrew, W. C. (2004). *The cultured chimpanzee: Reflections on cultural primatology*. Cambridge: Cambridge University Press.
- Millikan, R. (2006). *Varieties of meaning*. Cambridge: MIT Press.
- Möller, A., Pavlick, B., Hile, A. G., & Balda, R. P. (2001). Clark's nutcrackers *Nucifraga columbiana* remember the size of their cached seeds. *Ethology*, *107*, 451–461.
- Okrent, M. (2007). *Rational animals: The teleological roots of intentionality*. Athens: Ohio University Press.
- Rakoczy, H. (2015). Comparative metaphysics: The development of representing natural and normative regularities in human and non-human primates. *Phenomenology and the Cognitive Sciences*, *14*(4), 683–697.
- Ristau, C. (1991). Aspects of the cognitive ethology of an injury-feigning bird, the piping plover. In C. A. Ristau (Ed.), *Cognitive ethology: The minds of other animals* (pp. 91–126). New Jersey: Lawrence Erlbaum.
- Rowlands, M. (2015). Hard problems of intentionality. *Philosophia*, *43*(3), 741–746.
- Sachs, C. (2014). *Intentionality and the myths of the given: Between Pragmatism and Phenomenology*. New York: Routledge.
- Saidel, E. (2009). Attributing mental representations to animals. In R. Lurz (Ed.), *The philosophy of animal minds* (pp. 35–51). Cambridge: Cambridge University Press.
- Satne, G. (2016). A two-step theory of the evolution of human thinking. *Journal of Social Ontology*, *2*, 105–116.
- Schwitzgebel, E. (2010). Acting contrary to our professed beliefs or the gulf between occurrent judgment and dispositional belief. *Pacific Philosophical Quarterly*, *91*(4), 531–553.
- Searle, J. (2004). *Mind: A brief introduction*. New York: Oxford University Press.
- Sellars, W. (1981). Mental events. *Philosophical Studies*, *39*(4), 325–345.
- Slingloff, L., & Moore, R. (2018). Do chimpanzees conform to cultural norms? In K. Andrews & J. Beck (Eds.), *Routledge handbook for the philosophy of animal minds* (pp. 381–389). New York: Routledge.
- Sultanesco, O. (2015). Bridging the gap: A reply to Hutto and Satne. *Philosophia*, *43*(3), 639–649.
- Sultanesco, O., & Andrews, K. (2013). Are apes responses to pointing gestures intentional? *Humana Mente: Journal of Philosophical Studies*, *24*, 53–77.
- Tebbich, S., & Bshary, R. (2004). Cognitive abilities related to tool use in the woodpecker finch, *Cactospiza pallida*. *Animal Behavior*, *67*, 689–697.
- Vauclair, J. (1996). *Animal cognition: An introduction to modern comparative psychology*. Cambridge: Harvard University Press.
- von Rohr, R. C., Burkart, J. M., & van Schaik, C. P. (2011). Evolutionary precursors of social norms in chimpanzees: A new approach. *Biology and Philosophy*, *26*, 1–30.