# How to ascribe beliefs to animals

Tobias Starzak and Albert Newen, Ruhr-Universität Bochum

to appear in: *Mind and Language*

## 1. Introduction

The debate about animal belief has mainly been fueled by theoretical disagreement about the nature of belief (Andrews, 2016). According to a widely shared view, beliefs are representational states that express propositional content. Philosophers who subscribe to Fodor's *language of thought hypothesis* argue that this view is consistent with animal belief (Fodor, 1975; Cheney & Seyfarth, 2007; Tetzlaff & Ray, 2009). Others argue that the lack of language in animals shows that they have neither a language of thought nor beliefs (Stich, 1979; Davidson, 1982; Dennett, 1995). In a different vein, some philosophers argue that some animals have representational beliefs that cannot be understood as propositional (Camp, 2009; Glock, 2000; 2010) but as based on imagistic representational systems (like diagrams or maps) (e.g. Camp, 2007; Bermúdez 2003). On these accounts, animal belief is limited in various ways, e.g. it neither allows for second-order belief nor (because it lacks syntactic structure) for logical inference (Bermúdez 2006). Sometimes the question of animal belief is presented as depending on the question of animal concepts, such that those who deny animal concepts also deny the possibility of animal belief (Davidson, 1982; Stich, 1979). Other philosophers defend non-conceptual belief content (Bermúdez, 2003; Glock, 2010), or argue that animals can have concepts even in the absence of language (Newen & Bartels, 2007; Allen, 1999). These realist accounts of belief mentioned so far need to be distinguished from instrumentalist positions as introduced by Dennett (1995) that avoid any ontological commitment concerning animal mentality, and only demand for belief ascriptions to be justified that they allow successful predictions of behavior.

In this paper, we commit ourselves to a realist account of belief; our special contribution to the debate is a new suggestion for a minimal characterization of belief which can be outlined

as gradually unfolding into paradigmatic cases of linguistically expressible human beliefs, but which can be fruitfully applied to describe animal beliefs as well. We proceed as follows: *First* we present and discuss different versions of one of the most powerful arguments against animal beliefs. *Secondly*, we reject both premises of this argument in turn, and in doing so develop a notion of minimal belief. In the *final section*, we defend this notion by applying it to empirical cases of animal behavior.

Let´s start with the argument against animal beliefs. Roughly, it can be presented as follows:

*The argument against animal beliefs*

P(1) We are only justified in ascribing beliefs to animals, if we are justified in ascribing specific beliefs to animals.

P(2) We are not justified in ascribing specific beliefs to animals.

C We are not justified in ascribing beliefs to animals.

One version of the argument, spelled out by Davidson (1982), is best understood as an ontological argument: he argues that we have good reasons to think that animals do not have beliefs. Another version of the argument, presented by Stich (1979), is better characterized as a skeptical position that undermines our justification of belief ascriptions, although the idea of animal belief is consistent with the notion of belief he relies on.[1] Through criticizing both versions of the argument we will develop our own account of belief that allows us to account for animal beliefs as well. Following Sterelny (2003**)** we suggest that beliefs should not primarily be characterized as propositional, but more generally as informational states that are decoupled from specific immediate behavioral reactions. Furthermore, they are combinable with other informational states, and with a variety of different motivational states. This makes one and the same belief functionally relevant to a multitude of different behaviors (for a similar

[1] Other philosophers like Andrews (2016) or Carruthers (2004) have interpreted Stich's argument as an ontological argument as well. Section (2) of this paper is dedicated to defending our reading of the argument.

view see Carruthers, 2008)[2]. One consequence of this view is that states can be more or less belief-like, because both conceptual considerations and empirical observations support the view that decoupling and combinability come in degrees. On the one hand, this can be used to undermine the first premise of the argument against animal belief, since it makes a specific kind of behavioral flexibility a good indicator of belief-like representations, even if we cannot spell out the specific content of these beliefs (section 3). On the other hand, this account of belief links beliefs not only to flexibility in general, but to specific *behavioral and epistemic dispositions*, where epistemic dispositions need not be understood in terms of the *logical* inferences a subject with a set of beliefs will draw, but in terms of the more general ways subjects can combine information. In section 4 we will show how this view allows us to ascribe contents that are *approximate enough* to be explanatory, and we demonstrate using empirical examples that we are justified in ascribing beliefs to some animals.

## 2. The challenge: The argument(s) against animal belief

### *2.1 Premise one: Why ascribing content is necessary*

Intuitively, premise one seems innocuous. Since we do not have direct access to animal minds we have to infer the existence or absence of beliefs on the basis of the behavior an animal exhibits. Roughly, the inference from behavior to an underlying mental state is justified if ascribing this mental state explains the observed behavior better than alternative explanations in terms of non-mental states. However, in the case of belief, we have to make a distinction between *believing* something instead of *wanting* or *hoping* etc. something (i.e. the *attitude* a

[2] Without anticipating to much of the rest of this paper, despite some similarities to Carruthers´ view, the main difference between our view and his is that we take states to be more or less belief-like, depending on the degree to which they are decoupled from other states, while on Carruthers´ view beliefs do not allow for very varying degrees but are in need of being fully decoupled from motivational states. Instead of states being more or less belief-like, Carruthers takes differences in working memory to account for different degrees of behavioral flexibility found in different animal species (Carruthers 2013). We think that pointing to differences in working memory is an interesting suggestion that is compatible with our account, but from an evolutionary point of view, we take the either-or assumption - that states are either fully coupled (like Millikan´s *pushmi-pullyu states*) or fully decoupled - as less plausible a scenario than the idea of incremental decoupling we favor.

subject has towards a certain content), and the thing that is being believed (the *content*). Premise (1) rests on the assumption that an explanation of behavior in terms of the underlying attitude alone is not explanatory. Saying that a dog is barking up a tree because it has a belief state (without specifying what it is the dog believes) doesn't explain the dog's behavior. Since the main target of folk-psychological explanations in terms of beliefs and desires is to explain why a subject behaved one way rather than some other way those explanations rationalize specific behaviors in the light of *what is being believed* (in conjunction with what is being desired), i.e. we have to make the content explicit. For instance, we could explain the dog's behavior with its belief *that a squirrel is up in the tree* and its desire *to catch the squirrel*.

### *2.2 Premise two: Why we cannot ascribe content to animals*

While the first premise has a lot of intuitive plausibility, the second premise seems more debatable. Why wouldn't we be able to ascribe specific beliefs to animals? Two possibilities come to mind here. First, there could be evidence of absence, i.e. we might have reason to believe that animals don't have beliefs. In what follows we will refer to this as the *ontological* version of the argument. Secondly, we might simply lack evidence that animals have beliefs. In this case, the conclusion is somewhat weaker, stating that we do not – and will not ever – have reason to believe that animals have beliefs. We will refer to this version as the *epistemological version* of the argument. The arguments are very similar up to a certain point and we will consider both of them in turn.

#### *2.2.1 The ontological argument*

What evidence is there that animals do not have beliefs? There are countless examples from our everyday experience with animals that suggest otherwise. For instance, we think Malcolm's dog barks up the tree because *it believes* that a squirrel is in the tree (and *it wants* to catch it) (Malcolm, 1977). This explanation is analogous to how we make sense of human behavior in

our every-day language: in terms of what a subject *wants the world to be* (desires) and what it *believes it to be like*, respectively *what it believes to be an appropriate action to get what it wants*. In both cases the explanation rationalizes the behavior of a subject and allows us to successfully predict what these subjects are going to do next.

However, while it seems plausible at first to ascribe beliefs about *squirrels*, or *trees* to a dog, this initial plausibility gets lost once we consider what these ascriptions entail. Beliefs are thought to play the functional role they play for the production of behavior *in virtue of what they are about from the perspective of the subject*, i.e. in virtue of their intensional content (Davidson, 1982; Dretske, 1988). The intensional content of a belief, Davidson argues, is constituted by its inferential relations to all other beliefs a subject entertains, i.e. its specific place in a subject's *web of beliefs* (Davidson, 1982). For instance, the intensional content of my belief *that the squirrel is in the tree* is constituted by all my background beliefs about what squirrels are (little furry animals that live in woods, hide nuts etc.), what trees are (plants, made of wood, resource for paper, burn etc.), what it means to be *on* something, and so on. Since the content of those background beliefs itself depends on further background beliefs, Davidson argues, it only makes sense to ascribe beliefs to a subject if we are willing to accept that it has a rich net containing endless further beliefs:

> A belief is identified by its location in a pattern of beliefs; it is the pattern that determines the subject matter of the belief, what the belief is about. Before some object in, or aspect of, the world can become part of the subject matter of a belief (true or false), there must be endless true beliefs about that matter. (Davidson, 1975, p. 168)

Davidson argues that most of the beliefs within that pattern must be true, since "*false beliefs tend to undermine the identification of the subject matter; to undermine, therefore, the validity of a description of a belief as being about that subject*" (Davidson, 1975, p.168). In other words, if you have too many false beliefs about squirrels, at some point those beliefs stop being about squirrels, and start being about something else, or as Davidson puts it, "*too much mistake simply blurs the focus*" (Davidson, 1975, p.168).

Importantly, the point Davidson makes about the truth of most beliefs contained by any web of beliefs also holds for the *number* of beliefs needed to constitute content, or so he argues. Having too small a stock of beliefs blurs the focus too. Since the background beliefs concerning the objects a belief is about determine the intentional content of those beliefs, Davidson thinks that as long as we are not prepared to ascribe a vast amount of background beliefs to an animal (which even those sympathetic to the idea of animal belief reject), the claim that animals have mental states with intensional content is nonsensical. And since the intensional content fixes the reference, without intensional content animals cannot have *de re* content – and thus beliefs – either.

This conclusion is not identical with P(2) (We are not justified in ascribing specific beliefs to animals), but it should be uncontroversial that P(2) follows. Moreover, given that belief ascriptions entail a realist commitment to beliefs, the main conclusion of the argument follows, too:

*C(1): We are not justified in ascribing beliefs to animals.*

### *2.2.2 Stich's epistemological argument*

On the face of it, Stich´s epistemological version of this argument is very similar to Davidson's ontological argument. Like Davidson, Stich subscribes to belief holism and argues that this constitutes a problem for belief ascriptions to non-human animals. There are, however, two issues – one implicit and one explicit– on which Stich's position deviates from Davidson's. Both can be found in the following passage:

> Davidson and I differ only in that on my view content attribution presupposes a *shared* background of beliefs, while on his view what is presupposed is a *true* background of belief. The two views all but converge if we grant that most of *our* beliefs are true – a position few are inclined to deny. However, it is agreement not truth which is central, as is shown by the fact that if a subject simply *lacks* enough of our beliefs, though he holds no false ones, content attribution is undermined as surely as it would be if he held false beliefs. (Stich, 1979, p. 23)

The implicit disagreement concerns the nature of belief. The fact that Stich explicitly allows for the possibility that subjects have some, but significantly fewer beliefs than we do (Stich 1979, p.23, FN4), suggests that Stich rejects Davidson's claim that an endless number of beliefs is needed in order to have a single belief. Given this, animals *could* have beliefs although their web of beliefs is significantly less complex than our own.

However, the *explicit* disagreement between Stich and Davidson, concerning the requirements for belief *ascriptions*, shows why for Stich belief ascriptions are nevertheless problematic in the case of animals. Because the background of beliefs into which any given belief is embedded determines the content of a belief, Stich argues that we could not translate into our own conceptual repertoire the beliefs of other subjects unless we share the same background of beliefs. Given that this is not the case with animals, belief ascriptions are undermined for *epistemological* reasons. Wittgenstein's famous remark that even "[i]f a lion could speak, we could not understand him" (Wittgenstein, 2009, p. 235) captures the gist of Stich's argument: Even if an animal had beliefs, we could never say what it is they believe. It can be presented as follows:

| | |
|---|---|
| P(5) | We are only justified in ascribing specific beliefs to a subject if we have a shared background of beliefs. |
| P(6) | Humans have a web of beliefs containing an endless number of true beliefs. |
| P(4) | Non-human animals do not have a web of beliefs containing an endless number of true beliefs. |
| C(3) | We do not have a shared background of beliefs with non-human animals (follows from P(6-4).[3] |
| P(2)/C(4) | We are not justified in ascribing specific beliefs to animals (follows from P(5) and C(3)). |

[3] To be precise: C(3) follows from P6-4 given the plausible assumptions that substantial quantitative differences undermine the sharedness of a belief background.

Here is where we are: Davidson claims that there are good reasons to think that animals don't have beliefs, while Stich argues that there are no good reasons to think that they do. To reject Davidson's argument, the challenge is to show that having a single belief doesn't entail having an endless number of further beliefs. The challenge in rejecting Stich's argument is to either show that we can ascribe content to a creature whose belief system is considerably different from our own, or to show that we can be justified in ascribing beliefs to animals without ascribing specific beliefs. We will pursue both routes.

## 3 Against premise one: Explaining behavior without specifying content

### *3.1 Explaining behavior in context*

P(1) rests on the assumption that ascribing the ability to entertain beliefs without specifying the content of what is believed is not explanatory. Simply stating that Max has beliefs and desires does not explain why he took the umbrella because it does not rationalize Max's behavior in the light of *what* he believes (that it is raining) and desires (to not get wet).

The problem with this argument is that we slightly changed the subject. What we argued for is that only specific beliefs explain *specific* behavior (like Max's taking the umbrella), while the argument against animal belief depends on the claim that only specific beliefs explain behavior in a more general sense. This distinction might sound odd at first. What does it mean to explain *behavior in a more general sense*? The idea is that the behavior of subjects that have beliefs exhibits features that cannot be read off from single behaviors. In order to detect these features, we have to look at larger sets of behavior, or behavioral patterns. As we will argue in what follows, showing that a subject's behavioral patterns exhibit those features is both a necessary requirement for ascribing single beliefs and a sufficient requirement for ascribing the ability for having beliefs (or at least belief-like representations, see below).

### *3.2 Why we cannot ascribe beliefs on the basis of individual behaviors alone: beliefs and flexibility*

On the face of it, even the behavior of relatively simple animals often looks complex. For example, some ants carry dead conspecifics out of their nest (López-Riquelme et al., 2006). But do they have beliefs about their conspecifics being dead? Clearly not, because the observed behavior has been caused by some hard-wired stimulus-response mechanism.[4]

But why don't we explain hard-wired behavior in terms of beliefs and desires? One of the main reasons is that the predictive relationships between the states involved and the behavior that follows differ fundamentally in cases of hard-wired behavior on the one hand and cases usually explained in folk-psychological terms on the other. A core commitment of folk-psychology is that intentional behavior is the product of *two* kinds of states, informational states (beliefs), whose function it is to carry information about how the world is or how it could be; and motivational states (preferences, desires), whose function it is to rank possibilities of what the world could be like (Sterelny, 2003, p.6; for similar views see Papineau, 2003; Carruthers, 2004).[5] Importantly, beliefs can be combined with different further informational states as well as with different motivational states – and which behavior they cause crucially depends on which states they are combined with. As a result, there is no way of predicting a subject´s behavior on the basis of knowing that she has a certain belief A – unless we also know *what else* she believes, and *what she wants*. Ascribing the belief that the sun is shining can be part of an explanation for why Harry goes to the beach (because he *wants* to get a tan) but also for why he stays at home (because he wants to avoid skin cancer) (Kornblith, 2002). This is fundamentally different with rigid stimulus-response mechanisms. Here, merely knowing that

[4] In its most simple form this means that an organism's detection of some feature, here the smell of oleic acid, triggers the very same behavioral response every time.
[5] Searle (1983) characterizes the difference between beliefs and desires in terms of their *direction of fit*, mind-to world in the case of belief, and world-to-mind in the case of motivational states.

an ant has registered the smell of oleic acid is sufficient to predict more or less exactly what this ant is about to do next.[6]

According to many philosophers this predictive difference is a consequence of an architectural difference. Beliefs "are relevant to many possible actions, but functionally specific to none" (Sterelny, 2003, p. 50) and are thus best understood as *decoupled representations*. The strong one-to-one coupling between representational states and behavior in the case of hard-wired behavioral reactions on the other hand has been interpreted such that these representations – Millikan (1996) calls them pushmi-pullyu representations – have both informational and motivational (or directive) character at the same time.

Importantly, whether an observed pattern of behavior is part of an innate, rigid behavioral routine or the result of learning does not become transparent by simply observing one single behavior. To rule out hard-wired behavioral routines we have to check whether the animal is able to modify its behavior as a response to past experiences – whether it can learn.

However, ruling out innate behavioral routines is not enough to rule out pushmi-pullyu representations, as the case of the piping plover's *broken wing display* illustrates (Ristau, 1991). When a predator approaches on foot and the piping plover has a nest nearby, these birds fake a broken wing to present themselves as an easy target to protect their offspring. Once they have lured away the predator far enough from their nest, they take off and escape. But while these animals show some sophistication in their ability to learn who is a potential predator and who is not[7], the broken-wing display is the only reaction in their behavioral repertoire once something has been identified as a predator under these circumstances. Although their ability to learn introduces some flexibility to their behavior, this flexibility is restricted to *adjusting*

[6] This is true for the general behavior understood as *disposing of dead body from nest*. If we zoom in, however, there will be differences in how the ant executes this behavioral program, depending on the environmental specifics of the particular situation. But while ants need to show sensitivity and responsiveness to local peculiarities, this responsiveness is itself best understood in terms of triggered responses. We thank one of the reviewers for pointing that out.

[7] For instance, while they classify humans as potential predators, they can learn that human individuals which came close to - but never attacked - the nest before are not dangerous.

*the triggering conditions* for a fixed behavioral routine: their representation of predators is still tied to this specific behavior and thus best characterized as pushmi-pullyu representation as well.

Given that behavior caused by pushmi-pullyu representations can look sophisticated and is difficult to spot, we need a behavioral criterion for how to distinguish them from beliefs: we have to look at behavioral patterns and their development in order to determine whether the animal's representation is coupled to a single behavioral response, or whether it can be combined with other states to give rise to different and novel kinds of behaviors. Specifically, we have to check whether an animal can learn, and if so, whether it can use learned information about the world merely to adjust the triggering conditions for fixed behavioral routines, or whether it can use it for novel behaviors as well. Thus, the ascription of a specific belief is *never* justified unless we widen our focus from specific behaviors to a pattern of behaviors and its development.[8] This will be further spelled out as a profile of behavioral and epistemic dispositions.

### *3.3 Why ascribing specific beliefs is not necessary*

Following the argument of the last section the question *what an animal believes* doesn't even come up before we have good reason to think that it has beliefs. Hence, being able to classify a subject as a *believer* is not only possible without specifying the content of her beliefs – it is a precondition for it. Moreover, because beliefs have been characterized in a way that closely ties them to flexible behavior (involving novel actions), we can investigate whether an animal has beliefs by testing its ability to flexibly use information for different behaviors in different situations.

[8] A worry at this point is that content ascriptions may already be implicit in identifying patterns (e.g. identifying what is learned). To the extent that this is the case, a Dennettian intentional stance towards animals may be helpful. However, in contrast to Dennett, we merely treat this kind of content ascription as a heuristic to identify animals that merit closer inspection. In section 4 we spell out how we should investigate the content of animal belief-like states, realistically understood.

In the following we discuss a couple of possible objections. First, there's the worry that flexibility is too strong as a behavioral criterion and thus not necessary. On the other hand, one may also wonder whether it is too weak. Furthermore, one might argue that flexibility is too vague a concept to give us a good behavioral criterion. This last objection forces us to qualify our account to some degree.

### *3.3.1 Is flexibility necessary for decoupled representations?*

We characterized beliefs (in line with Sterelny, 2003) roughly as decoupled representations, i.e. as informational representations without motivational content that enable flexible behavior. However, if beliefs are just decoupled representations, it is not so clear whether behavioral flexibility follows, as the case of episodic-like memory illustrates. Episodic-like memory fulfills the criterion of decoupledness: information is being picked up and stored for later usage. The fact that subjects don´t act (more or less) immediately on this information supports the view that these states are essentially decoupled: they need to be combined with other states to trigger behavior. However, that does not entail that this information can be paired with a lot of further informational or motivational states: memory states could be restricted to drive a specific type of behavior only. So, if we identify beliefs with decoupled representations, flexibility does not seem necessary for having a belief.

We have two responses to this objection. *First*, even if flexibility was not necessary for having beliefs (informational decoupled representations), the fact that it is sufficient is all we need for our argument. If it is, and if we can empirically measure it, we have one reliable behavioral criterion for animal belief that is not touched by the content-problem.

*Secondly*, while it seems fair to point out the limits of flexibility in these birds, it is certainly wrong to claim that they lack flexibility altogether. Furthermore, this phenomenon of some, but limited, flexibility concerning the combinability of informational states with motivational states is what we find in many animals, and to varying degrees. In our view, the right perspective on

these cases is to not choose sides but instead emphasize both similarities and differences to the kind of behavior we account for in terms of full-blown beliefs, a view we further develop in section 3.3.3. For now, let´s simply call the states underlying limited flexibility *belief-like*."

### *3.3.2 Is flexibility sufficient for belief?*

The argument above characterizes beliefs as *informational states that are decoupled and inter-combinable, and whose content plays an explanatory role for the behavior of the subject of belief.* This characterization is consistent with beliefs having either determinate or vague content, as well as different formats of representation. A radical challenge for this view comes with the linguistic intuition, that beliefs have to have determinate fine-grained content and a peculiar representational format that accounts for compositionality, systematicity and productivity (Fodor, 1998). In this case, flexibility – the behavioral criterion for the (right kind of) informational character – would not be sufficient to ascribe belief-like representations.

There are a couple of things one can say in response to this. *First*, even if we want to anchor our notion of belief in the practice of ascribing beliefs to humans, it is doubtful that a restriction to linguistic belief is warranted. It is true that linguistic beliefs have quite determinate contents. But even in humans, it is the often implicit, coarse-grained beliefs rather than the explicitly expressed linguistic beliefs (with fine-grained contents) that guide action. As Beck (2012, p.535, our italics) writes:

> [I]nterpersonal ascriptions sometimes prove just as imprecise as interspecies ascriptions. And in fact, the same can be said of *intra*personal ascriptions. When each of us tries to linguistically characterize our *own* representations from core cognition, we are unable to find the right words. Indeterminacy is thus a problem not only of interpretation, but also of conceptualization. *When we linguistically conceptualize our own core representations, we distort their contents.* To borrow a Quinean maxim, indeterminacy begins at home.

Thus, linguistic belief is often something like a *post-hoc* rationalization that does not capture the exact content of our beliefs, but distorts it to some degree, either by translating a non-conceptual content into a conceptual one, or by translating content from one representational format to a different one. These cases are more the rule than the exception in humans and the reason we readily accept these cases as cases of belief is that the underlying states are essentially informational and re-combinable, and their content – even if we have only a rather rough grasp of it – rationalizes what they do.[9]

*Secondly,* even if this were not the case, we might want to abstract away from some of the features that are characteristic for human belief and focus on other features when we judge whether or not a subject has beliefs. After all, the explanatory role of belief content, the informational character, and the combinability of belief states seem to be particularly central features of belief according to folk-psychology, and if animals have states of that sort, they do share important features with humans. That is not to say that differences do not matter. If animals represent the world very differently in some other respect than humans do, e.g. in a different representational format, this is an interesting fact in itself. The problem reveals a dilemma that comparative questions concerning animal cognition often face when they only allow for "yes" or "no" as an answer: If all we are looking for are differences, we'll find them in one place or another. But this will often create the appearance of an unbridgeable gap between humans and non-human animals where there is none. If, on the other hand, we put too much emphasis on commonalities we run the risk of leveling things that are in fact very different in some respect. From a comparative point of view we have to find the right balance, and this requires more discriminating questions. We return to this in our answer to the third objection below.

[9] A similar idea can be found in Allen (2013).

*Thirdly*, as Allen & Bekoff (2007) argue, thinking about animal cognition can be a means to think about the nature of thought in general. Given this, it is not at all clear why we should anchor our understanding of belief in typical human performance. Rather, we should be open to the possibility that linguistic belief is a special case of a more fundamental ability. In the next section, we argue that this is in fact the case. Approaching the subject from this perspective, we have to be careful again to not jump too hastily to general conclusions concerning similarity or difference. But this provides us with a good reason to focus on some aspects of human belief and neglect others.

### *3.3.3 Is flexibility too vague?*

This brings us to the question of whether flexibility is too vague a concept to be a useful behavioral criterion. The problem is this: the degree of flexibility that the behavior of different kinds of organisms – frogs, beavers, cats, dogs, chimpanzees, New Caledonian crows, and humans – exhibit varies considerably. Since we'll already find some degree of flexibility in very simple animals, flexibility as a criterion for belief inflates the notion of belief up to a point where beliefs can no longer play a role in explaining behavioral differences. Why, then, should we care whether animals have it or not?

This objection calls for some qualification of our account. If we were to find out that frogs have beliefs, this would be thought big news; it should not be the result of a reconceptualization that makes it easy to meet the standard for having beliefs. However, the objection that flexibility as a criterion cannot account for a sharp line between belief and not-belief presupposes that there is such a sharp line. As we see it, this is not the case. Beliefs are decoupled and inter-combinable informational states. Decoupling as well as combinability are both a matter of degree (Sterelny, 2003). Representations can be functionally specific to more than one yet still a limited number of actions. Organisms may be able to learn how to use information for novel behavior, but this ability may still be constrained due to architectural features (such as

informational encapsulation) or limited computational resources. Furthermore, organisms can represent some aspects of their environment in a way that is less decoupled than other aspects. Vervet monkeys' representations of different kinds of predators (eagles, snakes, or leopards) are tightly coupled to specific, predator-dependent flight responses, but they can use information about social relations for many different actions (Cheney & Seyfarth, 1990). The more decoupled (and thus informational) representations become, and the more they can be combined with other states, the more *belief-like* they are.

The upshot is that differences in flexibility indeed reflect huge differences concerning the informational character of the underlying representational states. Flexibility, then, should not be seen as *the* criterion for belief, but rather as *a measure* for how belief-like an animal represents its environment.

Does that mean that, after all, animals do only have more-or-less belief-like states and not full-blown beliefs? Yes and no. As we argued above it is in many cases not obvious at all that folk-psychological explanations for humans are adequate, if we strictly reserve the label *belief* for linguistic belief. If we are fine with this practice in the case of humans, however, some animals can plausibly be said to have beliefs as well. If on the other hand we want to reserve the notion *belief* for states that are completely decoupled from any set of behavioral reactions, animals probably have merely belief-like states. However, even in humans having a belief is connected with typical dispositions to act although these are context-dependent and can be easily modified by learning. Thus, complete decoupling of belief from any behavioral dispositions is probably a misleading ideal.

But more importantly though, this story draws attention to the fact that any story that either suggests huge differences or no differences at all is not adequate. While animals may not have full-blown beliefs (with the level of decoupling of linguistic beliefs), some plausibly possess states that come very close: states that only differ in degree concerning features that are relevant

for beliefs; states for which psychological explanations (in terms of content) are adequate (given that we do find a way to ascribe content: see next section).[10]

Finally, thinking about belief-like representations as a gradual ability seems also to be compatible with the holism Davidson and Stich have in mind. If a holistic web of belief is too poor, for Davidson the right thing to say is that there are no beliefs at all. But which degree of complexity or of richness is needed? There does not seem to be a sharp line; there are only states that clearly count as belief and states that clearly don't. The same could be said using the notion of more-or-less belief-like states. But this suggestion offers a means to measure *how* belief-like representations are – by measuring the degree to which subjects can use information flexibly to guide their actions – and while Davidson is concerned with an either-or question, the story we offer promises to give a more detailed account of how (and in which respect) an animal's states differ or are similar to human belief states. Thus, while this vague notion of belief may sound unsatisfying for everyone who wants a simple yes-or-no answer to the question of animal belief, it turns out to be a feature rather than a bug in our account.

## 4. Against premise two: How to ascribe content to animals

In the previous section, we made a suggestion for how to find out whether an animal has belief-like representations while circumventing the content problem. While this undermines the general argument against animal belief (by challenging the first premise), it leaves open whether the second premise – that we cannot ascribe specific beliefs to animals – is true. Since this premise challenges much of the current practice in comparative psychology and cognitive ethology, we now turn to the question of whether (and in which cases) ascribing content to animal belief-states is explanatory and justified.

[10] For instance, the case of the piping plover illustrates how content ascriptions may capture some aspect of the underlying representation of an animal's environment and be explanatory, although the content will most likely be underspecified and the attitude towards that content not very belief-like.

First, we'll argue that the problem arises – at least in part – from an overly idealized conception of what constitutes acceptable belief-ascriptions in humans. Once we lower the requirements for adequate belief ascriptions – from literal *de dicto* ascriptions to approximate ascriptions – we see that the problem, although complicated, is not unsolvable. In a second step we propose a solution to this problem. We'll conclude that although we may not be able to grasp the exact content of animal beliefs, we can narrow it down enough for these ascriptions to be explanatory and justified.

### *4.1 Approximate ascriptions*

According to Stich, the problem with specifying the content of animal beliefs is that, strictly speaking, every ascription is false. Belief content is constituted by inferential relations to all the other beliefs that constitute a subject's holistic web of belief. Different inferential relations – due for instance to the considerably different quantity of beliefs any two belief-webs contain – constitute different contents. It is plausible that the web of beliefs of an animal differs in this respect. Thus, we could never translate animal beliefs into our own language – even if an animal had beliefs. As Carruthers (2004) puts it, we could never *co-think* an animal's thought. If we cannot say precisely what it is an animal believes, these ascriptions are not explanatory and thus should have no place in scientific disciplines such as comparative psychology or cognitive ethology.

A first thing to note here is that putting things this way somewhat overstates the problem with intensional/*de dicto* content. Consider the following example (taken from Schwitzgebel, 2015): Little Davy, after meeting his new teacher Ms. Sanchez, who's a Mexican, comes home crying from his first day at school. Later that day his mother tells his father that Davy thinks that Ms. Sanchez is too strict, or that the new Mexican teacher is too strict. These referentially transparent (*de re*) ascriptions appear to be fine and, more importantly, they *are* explanatory of Davy's crying, even if Davy neither knows that Ms. Sanchez is Mexican nor that her name is

Ms. Sanchez. If this is so, they are explanatory although they fall short of the requirement to capture the intensional content of Davy's thought. According to Schwitzgebel (2015), and contrary to Stich and Davidson, "the standard view takes belief attributing sentences to be systematically ambiguous between a referentially opaque, *de dicto* structure and a referentially transparent, *de re* structure". One reason for this may be that our ability to explain behavior in terms of beliefs and desires would otherwise be severely limited, even in the case of humans. Taking the example above, which terms we use in everyday practice to report a belief not only depends on how the believer represents the world, but also on the background knowledge of the third person to whom we are reporting this belief (see Ueda, 2015). If Tom thinks that Jake stole Connor's car, but Connor knows Jake only under his street-name T-Bone, I had better report to Connor that Tom thinks T-Bone stole his car if I want him to understand whom Tom thinks stole his car.

Now it is certainly true that not all belief ascriptions are intended to communicate something to a third party but in many cases to make sense of the behavior of others for ourselves. But if one embraces holism, belief ascriptions to other humans appear to be always approximates – and not just for pragmatic reasons (Schwitzgebel, 2015). My holistic web of beliefs may exhibit a higher degree of isomorphism to that of another human than to that of a dog or a chimpanzee, but it will also differ in many respects to that of other humans. My own web of beliefs will be considerably different to that of members of very different cultures, like the Haroli from Papua New Guinea, or to that of people who lived 2000 years ago, who lack many of my background beliefs that are dependent on my own culture or on more recent scientific discoveries. Yet it doesn't seem problematic at all to explain their behavior in terms of beliefs and desires. Furthermore, if the content of my beliefs is determined by the inferential relations to all my other beliefs, the problem easily generalizes to all humans, including members of my own culture living at the same time as me, since no two webs of belief are identical. Thus, even in the case of humans, isomorphism understood as identity between two webs of belief is clearly

out of reach.[11] The fact that we still accept the practice of explaining behavior in terms of beliefs as sound shows that what matters is the degree of similarity rather than identity. One may argue that any claim about which degree suffices will have an air of arbitrariness.

Indeterminacy seems like a shortcoming when we take the beliefs we want to report to have fine-grained contents. But as we argued in section 3.3.2, even in humans it is often implicit and coarse-grained beliefs rather than explicit (linguistic) beliefs (with fine-grained contents) that guide action. For non-linguistically manifest beliefs, the problem of precisely capturing their content in words arises not only from a lack of knowledge about how exactly a subject represents the world, but also from translating a state from a non-linguistic format to a linguistic one. As Quine put it, "radical translation [and thus: indeterminacy] begins at home" (Quine, 1968, see also section 3.3.2 of this paper). Although the linguistic content of a belief is fine-grained it is always just a rough characterization of a complex state of mind which we use *to mark sufficient similarity of contextually relevant behavioral and epistemic dispositions*. Hence, we argue that we are not committed to a dangerous arbitrariness but just to a sensitivity of belief ascriptions to a pragmatic context, according to which we normally can easily agree about the relevant aspects of similarity. Context-sensitivity is not a disastrous feature of beliefs and belief ascriptions at all; indeed, it is widely accepted (Recanati, 2000).

To sum up: Even if it turns out that we are not able to literally translate animal thoughts into natural language, these ascriptions do not necessarily cease to be explanatory. For humans, we accept approximate ascriptions all the time. In the next section, we´ll have a closer look at the conditions under which this works for animals as well.

[11] It should be noted that Stich does not think of isomorphism as an identity criterion, but rather as one of sufficient similarity. The problem is, however, that he presupposes that we have a clear grasp of which degree of similarity suffices.

### *4.2 Animal beliefs with specific contents*

When are we justified in language-based belief ascription? We suggest that three criteria have to be satisfied. *First,* the cognitive system must exhibit sufficiently complex behavior, as spelled out above. This can be characterized by a certain cluster of *behavioral dispositions*. *Secondly*, as a demand concerning the *structural organization*, activities belonging to this behavioral cluster are based on internal states that are largely informational and inter-combinable. *Thirdly*, these informational states are connected with *epistemic dispositions* which involves *sensitivity* for new information, an *ability to cluster or categorize new information* and a demand for *adjustment of the relevant informational states* in the light of new evidence. These minimal epistemic dispositions include rudimentary memory abilities which enable a cognitive system to activate stored information and learning abilities to cluster or integrate new information.[12] These three criteria are different but not independent features of beliefs. Complex behavioral dispositions are realized by a certain structural organization of the relevant mental states. Furthermore, since the minimal epistemic dispositions include the ability to cluster or categorize new information about the same type of entity, this typically involves the ability to recognize the same object, property and type of situation. The latter means that beliefs typically include *structured information* about these features that involve different types of basic components like events, substances, directions, or maps depending on the way a cognitive system experiences and acts in the world.[13] This overlaps with the demands for the structural organization of beliefs and it enables new re-combinations of worldly information, e.g. about objects and their changing properties.

[12] This may invite the question which role memory abilities play for having beliefs. Since we argue that minimal flexible behavior is based on sensitivity for new evidence and learning, we presuppose the ability to memorize new evidence and to integrate it into already stored relevant information. This presupposes learning abilities and a memory capacity but does not make very specific presuppositions concerning the architecture of the memory system. Some constraints come with the condition of a structural organization of informational states such that components can be recombined in different ways. Thus, the memory system must be able to store informational components and the learning strategies must be able to handle them.

[13] Components of the informational system of bees, for instance, are substances (e.g. nectar source), directions (or angles), and the position of the sun. See also FN 16.

The third criterion needs to be spelled out in more detail. The informational states typically include a certain way of representing objects, including that objects are registered as permanent entities with different properties. Relevant properties of objects include the functions of an object (where these are registered as affordances) as well as shape, color, material etc. An important condition for such an informational representation to be a candidate to be a belief is that an object in one visual scenario can be represented as having different properties, and one and the same property can be represented as being instantiated in different objects. Irene Pepperberg's famous grey parrot Alex fulfilled this requirement: he could represent an object as red (in contrast to four other colors), as round (in contrast to three other shapes), and as wooden (in contrast to two other materials) (Pepperberg, 1999). This exemplifies a structured representation of objects, their (relational) affordances and (further) properties. This minimal structuring enables *systematic associations of an object and a property* and an easy *re-coupling of the associations* which constitute a predictable cluster of behavioral and epistemic dispositions. We have to presuppose this also in the case of understanding the causal features of objects and their creative use, e.g. when chimpanzees use water to float a peanut (Tennie, Call & Tomasello, 2010), or when ravens use stones to raise the level of water to reach swimming food (Jelbert et al., 2014), or are able to solve a sequence of causal tasks (Taylor et al., 2007).

These minimal epistemic dispositions can develop gradually into more and more sophisticated dimensions, which then leads to increased complexity of all three epistemic features: higher sensitivity to worldly information, more fine-grained representations of entities in the world and a more sophisticated integration of worldly information into a beliefs system. The combination of these features works as follows: having beliefs comes with the epistemic demands to adjust them in the light of new evidence while this is realized in accordance with the principle of correspondence (account for the reality) and the principle of coherence (keep

the information about the world consistent (at least contextually), keep it simple and as interconnected as possible).

Together, these abilities form a cluster of behavioral and epistemic dispositions that justifies the ascription of beliefs, because these ascriptions enable fruitful explanations and predictions of the animal behavior as well as provide a model of the cognitive organization of the worldly information and of the dynamics of adjusting an informational system.[14] Evolutionary and developmentally, basic beliefs can unfold through a systematic enrichment of the epistemic dispositions into full-blown fine-grained beliefs (like linguistic beliefs). A rough characterization of this development can be made by highlighting that – picking up the intuition of holism – basic beliefs have a low level of systematic interconnections and coherence, relying on cognitive heuristics which are fruitful in a certain environment, while language-based beliefs can be characterized as a very rich network of inferential or logical relations between beliefs which at the high end of the spectrum become increasingly domain-independent.

An example should help to illustrate the adequacy and the explanatory value of attributing basic beliefs. Trained in an eight-arm maze rats are able to register three daylight situations (morning, midday, evening), to distinguish different types of food (*normal* food and chocolate), and of course to register the spatial organization of the maze. Rats can learn to understand that *if they find normal food in arm2 in the morning, they will find chocolate in arm7 at midday*. On our account, their behavioral pattern and the organizational structure of the informational process justifies the following belief ascription: given the evidence of *finding normal food in arm2 in the morning, I will find chocolate in arm7 at midday*. Why is this not better described as procedural knowledge or *knowing how*, which – as many agree – we can account for without presupposing beliefs (Jung & Newen, 2010; Noë, 2005)? The answer is that it involves an

[14] We leave it open whether basic beliefs are already conceptual or whether we have to allow for non-conceptual beliefs, since there is no consensus on the conditions of possessing concepts. Our characterization is consistent with allowing for non-conceptual basic beliefs, while concept possession needs some additional epistemic abilities such as those spelled out in Newen & Bartels, 2007.

informational state with a (partially) decoupled representation of the type of food that can be combined with representations carrying information about different day-times and different locations. Earlier studies have shown that rats (at least a large group of them) can do spatial orientation based on landmarks (realized by the hippocampus), in contrast to those groups of rats which rely on rigid procedures (realized by basal ganglia) (Packard & McGaugh, 1996). The informational state of rats which have learned to behave in the maze according to a conditional is best characterized as *structured into components* of <object-type; location; time> and those informational states can quickly be reorganized by either learning new conditionals or, even more demanding, by learning to understand the eight-arm maze to have been newly restructured (e.g. by learning that some arms can be ignored in certain task situations, Wilson et al., 2013). The way different informational components can be combined is rather flexible and thus comes with an interesting degree of epistemic dispositions: the rats are sensitive to the findings at each time point, they update their expectation according to these findings, and they can minimally integrate their worldly information as can be shown by the fact that the same eight arm maze can be informationally represented as realizing two independent systems while each consists of three different relevant arms (and two arms remain irrelevant) (Crystal et al., 2013). If we were to test children in an eight-arm maze under the same conditions, it would depend on the age whether the epistemic conditions would remain almost at the same level (which is probably the case for 1-year-olds)[15] or would substantially go beyond. It is clear that 5-year-olds develop a significantly increased understanding of the full conditional that *if they find normal food in arm2 in the morning, they will find chocolate in arm7 at midday.* 5-year-olds have the additional cognitive capacities of second-person-perspective-taking and of passing the explicit false-belief task. Thus, the same structured components <object-type;

[15] We know from the systematic comparison of chimpanzees and orangutans with children of age 2.5 years that in non-linguistic physical tasks all three groups perform equally well, while the children are only better in non-linguistic social tasks (Herrmann et al. 2007). We do not have such a systematic comparison for rats but we can expect similar results even though the age comparison may be somewhat different.

location; time> are connected into a much richer network of similar structured components, realized as language-based beliefs. As adults, with many further epistemic abilities, the understanding gets even richer, marked by an explosion of the inferential relations between relevant informational states.[16]

Given these differences, why, then, are belief ascriptions nevertheless explanatorily helpful and thus justified in all three cases? Because there is a proven overlap of core cognitive processes shared by humans and animals, such that right from the start there is a sufficient amount of behavioral and epistemic abilities according to our three criteria which humans share with many other species. Although many evolutionary basic abilities are innate, they are often connected with core cognition including representations of objects, space, number, agency and social relations (Kinzler & Spelke, 2007; Carey, 2009). These shared cognitive abilities are the basis for the systematic unfolding of complex cognition during infancy. In some cases, we have very detailed descriptions of shared abilities and their development, e.g. the representation of number (Carey, 2009). Object permanence tests show that infants as well as many animals develop representations of objects as stable (Baillargeon & DeVos, 1991). Moreover, many animals – such as birds, apes and some types of dog – have been proven to develop an understanding of object-property (Pepperberg, 1999) or of object-action structure (Pilley & Reid, 2011). Dogs, e.g. can learn to associate a large number of acoustic symbols with individual objects and understand structured orders (like "bring Timi", where Timi is one out of 200 objects known by a dog (Kaminski et al., 2004)). Furthermore, strong evidence has been found of an understanding of the causal role of objects (Taylor et al., 2009; Seed et al., 2011).

[16] The famous waggle dance of bees to communicate amount and location of a food source provides us with a nice example for rather simple behavior that can still be adequately characterized in terms of belief, given our gradual characterization. Bees are equipped with a built-in solar ephemeris function, but they still have to *learn from their experience* with the sun's changing position over the horizon at different times of the day the locally appropriate values for the parameters of the solar ephemeris function. Furthermore, a bee has *to store and update* the information about the relation between the direction of the source from the hive and the angle of the sun, since this information was usually acquired many minutes earlier than the waggle dance is realized (Menzel & Eckoldt, 2016). Thus, while bees cannot use information as flexible as scrub-jays or rats, bee cognition does involve both learning and some structured representation of location, direction, and richness of the source.

These discoveries strongly indicate that we share with many animals important aspects of our behavioral and epistemic dispositions in dealing with objects: we are able to recognize them, we take them to be permanently existing and as having different properties which are relevant for different challenges, we categorize them, expect them to play relevant causal roles in certain situations and integrate this information, etc. These commonalities give us a sufficient basis for ascribing beliefs about objects to both infants and many animals that fulfill these criteria.

How can we account for obvious differences? The differences do not block the attribution of a belief, but they are accounted for by the fact that we understand belief attributions always *as implicitly indexical*, relative to a behavioral and epistemic profile of a relevant group or species: When we say that a three-year-old child believes that uncle Bob will come to visit tomorrow, we know even at a folk-psychological level that this belief is characterized relative to the epistemic abilities of infants and not relative to those of adults. In other words, we know that 3-year-olds have informational states of this type as separate from desires (since they handle desires as quite decoupled; Wellman, 2014), and it is very fruitful to describe the informational state as structured: being about a person, the property of coming and the time of the next day. It deserves to be called a belief since it explains why the child starts to paint a picture for uncle Bob, whereas last time she prepared a special soft drink for him. Nevertheless, the child has no understanding of the definition of the concept UNCLE. When understanding a belief ascription to a 3-year-old, we implicitly account for her general cognitive profile and describe a cluster of behavioral and epistemic dispositions relative to this. Thus, a further crucial benefit of the way we propose to characterize beliefs is that we can both capture our everyday practice as well as defend common scientific practice in cognitive ethology and comparative psychology.

**Conclusion**

We have discussed and rejected one of the most prominent arguments against animal belief. In doing so, we developed a gradualist account of belief where beliefs can be characterized as follows: For a cognitive system to have a belief means that the system has (i) *informational states that are sufficiently decoupled from motivational states and that are (to varying degrees) inter-combinable with other informational states as well as with different motivational states.* These informational states have (ii) *a minimal structural organization, typically (but not exclusively) involving elements that represent objects, properties or substances*. Moreover, these states are (iii) *connected with minimal epistemic dispositions* including (iiia) *a sensitivity for new information,* (iiib) *an ability to cluster or categorize new information when it concerns the same situation type or the same object, etc.* and (iiic) *the ability to adjust the relevant informational states in the light of new evidence*.

On the one hand, this account makes the question whether animals have beliefs empirically tractable, even in cases in which we are not able to spell out *what* an animal believes: the degree to which relevant representations are informational and inter-combinable (and thus: belief-like) can be measured in terms of behavioral flexibility which an animal exhibits. On the other hand, we also indicated a solution to the problem of content ascription. We have shown that the skeptical content problem is rooted in an over-intellectualistic picture of the nature of belief and belief-ascriptions in humans. Since beliefs are linked to behavioral and epistemic dispositions and we share some of these dispositions with many animals via our core cognition, we have enough common ground to spell out the content of animal beliefs closely enough for these ascriptions to be explanatory and allow predictions of behavior and cognitive processing. Discussing empirical examples, we have shown that belief ascriptions are useful for many animals including scrub-jays, rats, dogs, or corvids. Furthermore, our account is suited to investigate more specific questions concerning cognitive similarities between humans and non-human animals.

*Acknowledgment*

We are especially grateful to Hanjo Glock, who helped to initiate the paper, for an intense collaboration. We´d also like to thank Frances Egan, Robert Matthews, Bence Nanay, the Research Group of Albert Newen, and two anonymous reviewers for their constructive feedback that helped to improve the paper a lot.

This article was supported by the following projects/gefördert durch die Deutsche Forschungsgemeinschaft (DFG) (NE 576/14-1; SCHL 588/3-1 "Understanding Actions and Reasons") and Mercator Research Center Ruhr (Pr-2016-0016: "Automaticity in Thought an Action"). We are grateful for this support.

## *References*

Allen, C. (1999). Animal Concepts Revisited: The Use of Self-Monitoring as an Empirical Approach. *Erkenntnis* 51(1), 33–40.

— & Bekoff, M. (2007). Animal Minds, Cognitive Ethology, and Ethics. *Journal of Ethics* 11, 299–317.

— (2013). The Geometry of Partial Understanding. *American Philosophical Quarterly*, 50, 3, 249–262.

Andrews, K. (2016). Animal Cognition. *The Stanford Encyclopedia of Philosophy* (Summer 2016 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2016/entries/cognition-animal/>.

Baillargeon, R. & DeVos., J. (1991). Object Permanence in Young Infants: Further Evidence. *Child Development* 62(6), 1227-1246.

Beck, J. (2012). Why we can´t say what animals think. *Philosophical Psychology* 26(4), 520-546.

Bermúdez, J. L. (2003). *Thinking Without Words*. Cambridge, MA: MIT Press.

— (2006). Animal Reasoning and Proto-Logic. In S. Hurley and M. Nudds (Eds.), *Rational Animals?* Oxford: Oxford University Press, 127–138.

Camp, E. (2009). A Language of Baboon Thought? In R. Lurz (Ed.), *Philosophy of Animal Minds*. New York: Cambridge University Press.

— (2007). Thinking with Maps. *Philosophical Perspectives* 21, 145–182.

Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.

Carruthers, P. (2004). On Being Simple Minded. *American Philosophical Quarterly*, 41(3), 205-220.

— (2008). Meta-cognition in Animals: A Skeptical Look. *Mind and Language* 23(1), 58–89.

— (2013). Animal minds are real, (distinctively) human minds are not. *American Philosophical Quarterly*, 50(3), 233-248.

Cheney, D. L. & Seyfarth, R. M. (1990). *How Monkeys See the World: Inside the Mind of Another Species*. Chicago: The University of Chicago Press.

— (2007). *Baboon Metaphysics: The Evolution of a Social Mind*. Chicago: University of Chicago Press.

Crystal, J. D., Alford, W.T., Zhou, W. & Hohmann, A.G. (2013). Source Memory in the Rat. *Current Biology*, 23, 387-391.

Davidson, D. (1975). Thought and Talk. In D. Davidson, *Inquiries into Truth and Interpretation*, 155-170, Oxford: Oxford University Press.

— (1982). Rational Animals. *Dialectica*, 36(4), 317-327.

Dennett, D. (1995). Do Animals Have Beliefs? In H. Roitblat and J. Meyer (eds.), *Comparative Approaches to Cognitive Science*. Cambridge, MA: MIT Press, 111–118.

Dretske, F. (1988). *Explaining Behavior. Reasons in a World of Causes*. Cambridge, MA: MIT Press.

Fodor, J. A. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.

— (1998). *Concepts. Where Cognitive Science Went Wrong.* Clarendon Press.

Glock, H.-J. (2000). Animals, Thoughts, and Concepts. *Synthese* 123, 35–64.

— (2010). Can Animals Judge? *Dialectica* 64, 11–33.

Hanus, D., Mendes, N., Tennie, C. & Call, J. (2011). Comparing the Performances of Apes (*Gorilla gorilla*, *Pan troglodytes*, *Pongo pygmaeus*) and Human Children *(Homo sapiens)* in the Floating Peanut Task. *PLoS ONE* 6(6), 1-13.

Herrmann, E., Call, J., Hernández-Lloreda, M. V., Hare, B. & Tomasello, M. (2007). Humans Have Evolved Specialized Skills of Social Cognition: The Cultural Intelligence Hypothesis. *Science* 317, 1360-1366.

Jelbert, S. A., Taylor, A. H., Cheke, L. G., Clayton, N. S. & Gray, R. (2014). Using the Aesop´s Fable Paradigm to Investigate Causal Understanding of Water Displacement by New Caledonian Crows. *PLoS ONE* 9(3), 1-9.

Jung, E.-M. & Newen, A. (2010). Knowledge and Abilities: The Need for a New Understanding of Knowing-how. *Phenomenology and Cognitive Sciences* 9, 113-131.

Kaminski, J., Call, J. & Fischer, J. (2004). Word Learning in a Domestic Dog: Evidence for "Fast Mapping". *Science* 304, 1682-1683.

Kinzler, K. D. & Spelke, E. S. (2007). Core Systems in Human Cognition. In C. von Hofsten and K. Rosander (eds.), *Progress in Brain Research*, Vol. 164 (14), 257-264.

Kornblith, H. (2002). *Knowledge and its Place in Nature*. Oxford: Oxford University Press.

López-Riquelme, G. O., Malo, E.A., Cruz-López, L. & Fanjul-Moles. M. L. (2006). Antennal Olfactory Sensitivity in Response to Task-related Odours of Three Castes of the Ant Atta Mexicana (Hymenoptera: formicidae). *Physiological Entomology* 31, 353–360.

Malcolm, N. (1977). Thoughtless Brutes. In N. Malcolm, *Thought and Knowledge: Essays by Norman Malcolm*. Cornell University Press.

Menzel, R. & Eckoldt, M. (2016). *Die Intelligenz der Bienen. Wie sie denken, planen, fühlen und was wir daraus lernen können.* Knaus.

Millikan, R. (1996). Pushmi-Pullyu Representations. *Philosophical Perspectives* 9, 185–200.

Newen, A. & Bartels, A. (2007). Animal Minds and the Possession of Concepts. *Philosophical Psychology* 20(3), 283–308.

Noë, A. (2005). Against Intellectualism. *Analysis* 65(4), 278-290.

Packard, M. G. & McGaugh, J. L. (1996). Inactivation of Hippocampus or Caudate Nucleus with Lidocaine Differentially Affects Expression of Place and Response Learning. *Neurobiology of Learning and Memory* 65, 65-72.

Papineau, D. (2003). The Evolution of Means-End Reasoning. In D. Papineau, *The Roots of Reason. Philosophical Essays on Rationality, Evolution and Probability.* Oxford: Clarendon, 83–129.

Pepperberg, I. M. (1999). *The Alex Studies.* Cambridge, MA: Harvard University Press.

Pilley, J. W. & Reid, A. K. (2011). Border collie comprehends object names as verbal referents. *Behavioural Processes* 86, 184–195.

Quine, W. V. O. (1968). Ontological Relativity. *Journal of Philosophy* 65, 185-212.

Recanati, F. (2000). *Oratio Obliqua, Oratio Recta. An Essay on Metarepresentation*. Cambridge: MIT Press.

Ristau, C. A. (1991). Aspects of the Cognitive Ethology of an Injury-feigning Bird, the Piping Plover. In C.A. Ristau and D.R. Griffin (Eds.), *Cognitive Ethology: The Minds of Other Animals: Essays in Honor of Donald R. Griffin.* Hillsdale: Erlbaum, 91–126.

Schwitzgebel, E. (2015). Belief. *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2015/entries/belief/>.

Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind.* Cambridge University Press.

Seed, A., Hanus, D. & Call, J. (2011). Causal Knowledge in Corvids, Primates, and Children. In T. McCormack, C. Hoerl and S. Butterfill (eds.), *Tool Use and Causal Cognition*. Oxford: Oxford University Press, 89-108.

Sterelny, K. (2003). *Thought in a Hostile World. The Evolution of Human Cognition.* Oxford: Blackwell.

Stich, S. P. (1979). Do Animals Have Beliefs? *Australasian Journal of Philosophy* 57(1), 15-28.

Tennie, C., Call, J. & Tomasello, M. (2010). Evidence for Emulation in Chimpanzees in Social Settings Using the Floating Peanut Task. *PLoS ONE* 5(5): e10544.

Taylor, A., Harwood, Hunt, G. R., Holzhaider, J. C. & Gray, R. D. (2007). Spontaneous Metatool Use by New Caledonian Crows. *Current Biology* 17, 1504-1507.

—, Roberts, R., Hunt, G. & Gray, R. (2009). Causal Reasoning in New Caledonian Crows. *Communicative and Integrative Biology* 2(4), 311-312.

Tetzlaff, M. & Rey, G. (2009). Systematicity and Intentional Realism in Honeybee Navigation. In R.W. Lurz (ed.), *The Philosophy of Animal Minds*. Cambridge: Cambridge University Press, 72-88.

Ueda, T. (2013). *Telling What She Thinks*. DeGruyter.

Wellman, H. M. (2014). *Making Minds: How Theory of Mind Develops*. Oxford: Oxford University Press.

Wilson, A. G., Pizzo, M.J. & Crystal, J.D. (2013). Event-based Prospective Memory in the Rat. *Current Biology* 23, 1089-1093.

Wittgenstein, L. (2009). *Philosophical Investigations*. Wiley-Blackwell.