CHAPTER 33

EVOLUTIONARY PRECURSORS OF NEGATION IN NON-HUMAN REASONING

MANUEL BOHN, JOSEP CALL, AND CHRISTOPH J. VÖLTER

33.1. Introduction

In the philosophical literature, negation has been firmly tied to natural language or equivalent symbolic systems (e.g. Horn 2001a; Bermúdez 2003). So far, there is no evidence that any nonhuman animal species might possess a communication system that relies on symbolic representations. As a consequence, animals might be 'out of the game' when it comes to negation. In this chapter we want to discuss findings from animal cognition research that suggest that animals think and reason in ways that resemble reasoning based on negation. In the main part of the chapter, we offer a more detailed review of some particularly promising findings from the literature on individual and social reasoning. While intriguing, we conclude that none of these provides definite evidence for reasoning based on propositional negation. In the third part of the chapter, we suggest ways to approach the question of propositional negation in animals empirically. Finally, we sketch out some ideas about the role of negation in the evolution of propositional thought more generally.

33.2. Facets of negation

In this chapter, we follow Bermúdez (2003, 2006) and distinguish between proto-negation and negation proper (hereafter negation). The core of this distinction is that while proto-negation operates on at least two discrete representations that are contraries, negation

operates on a single representation (a proposition) that is evaluated as being the case or not. In Fregean terms (1892), negation operates on intensions, that is on propositions that represent the specific aspects of the world. Negating such propositions implies that whatever the intensional aspect is, is not the case. Proto-negation, on the other hand, operates on mutually exclusive extensions, that is, representations of certain states of affairs in the world. These representations can be thought of, to some degree, as pictorial representations of a given scene. As a consequence, a certain state of affairs is negated by finding out that the other member of the contrary pair is the case. Proto-negation nevertheless implies structured thoughts with elements that may reappear in different thoughts.

33.2.1. Proto-negation

Proto-negation is thought to be based on contrary pairs, such as 'presence–absence' or 'danger–safety'. Let us illustrate the distinction between proto-negation and negation with an example. In a study by Call (2004), great apes were faced with two potential hiding locations of a food reward. In one condition, they were shown that one of the hiding locations, say location A, was empty. When making a choice, apes chose location B and therefore presumably reasoned that if the reward is not in A, it must be in B. On a proto-negation account, rather than negating the presence of a desired food reward in A, individuals might have arrived at the conclusion by reasoning based on mutually exclusive states, namely that the reward is either present (absent) in A or present (absent) in B. When seeing that the reward is absent from A, apes concluded that the reward is present in B. The consequence of negating one state of affairs is thinking that a different state of affairs is the case.

Let us elaborate: The straightforward way to express this inference in a formal way would be the disjunctive syllogism. 'P or Q', 'not-P' therefore 'Q', with P being 'food present in A' and Q being 'food present in B'. This formulation, however, implies the ability to negate the proposition expressed by P, which should not be possible on a proto-negation account. Bermúdez (2003) describes proto-negation as a form of predicate negation. Instead of negating the entire proposition P, only the predicate 'present' is negated. A negative predicate, he argues, can be substituted by a different predicate that has the same extension. In the above example, it would mean concluding 'absent' instead of 'not present'. Both describe the same situation but absent is not constructed from present and therefore does not require an intension-preserving representation of present. As a consequence, proto-negation as an operation allows an individual to switch between the elements of a contrary pair.

Reasoning based on proto-negation therefore requires discrete contrary pairs (or sets, each with a limited number of elements). This begs the question of how animals come to form contrary pairs. Contrary pairs need not be mutually exclusive from a logical perspective but they should be mutually exclusive based on the experience of a given individual. If, for example, an animal experiences that two sorts of fruit are never available at the same time, the animal might come to form a contrary pair allowing it to infer the absence of one kind of food based on seeing the other kind. However, contrary pairs might be formed inferentially based on knowledge of spatio-temporal or object–object relations (Völter and Call 2017). For example, if an animal has some understanding of object solidity, it might

come to form the contrary pair that a board lying on the ground and an object being hidden under the board are mutually exclusive (see e.g. Call 2004). The way that contrary pairs are formed determines the range and flexibility of how an individual may be able to use proto-negation, but it does not necessarily alter the reasoning process itself.

To sum up, the important distinction is that conclusions based on proto-negation are limited to discrete states of affairs depending on the number and types of contrary pairs an individual is sensitive to. Which contrary pairs an individual is sensitive to depends to a large part on direct experience.

33.2.2. Negation

Whereas proto-negation operates based on contrary pairs, negation, on the other hand, operates on negating intensional aspects of a given state of affairs. As a consequence, it allows individuals to conclude a wider range of potential alternatives. In our example above, from the fact that the food is not in A could also follow that the experimenter ate the reward or that an invisible goblin stole it. These alternatives are certainly less plausible, but nevertheless possible. Negation only implies that one state of affairs is not the case. From a representational point of view, being able to negate a state of affairs requires an individual to form a representation that captures the intension of a proposition with which it could be described. To illustrate this point, think of an animal that represents states of affairs in a purely pictorial way. Such an individual would not be able to differentiate between two distinct states of affairs that have the same pictorial extension, say an empty box. For example, to differentiate between the thoughts that the banana is not in the box and that the grape is not in the box, they would need to be structured in a way that they capture a certain aspect (or perspective) on a visual scene. Among others, Bermúdez (2003, 2006) argues that the necessary structure is only provided by representations that are based on a system analogous to human language. By convention, symbols allow the representation of intensional aspects of experience. This is argued to be a necessary prerequisite to negate this very aspect. We do not necessarily commit to the claim that natural language is a necessary prerequisite for negation, but we do acknowledge that thoughts need to be structured so that they allow for representing intensional aspects.

33.2.3. What do you need negation for?

In the following, we explore the evolutionary relevance of negation by asking how the capacity for negation might impact an organism's cognition and as a consequence its behavior. In each case, we start with considering empirical findings in animals that seem to be related to negation. We discuss what animals are capable of doing and if/how negation proper would potentially alter this ability.

33.2.3.1. *Individual reasoning*

In their natural environment animals often face negative information. Negative information refers here to the perceptual absence of a stimulus or cue, that is a situation in which an

individual currently cannot perceive a stimulus. Perceptual absence includes occlusion events and situations in which there is actual evidence about the absence of a stimulus (henceforth: explicit absence). A central question is whether animals can distinguish different types of negative information, that is between occlusion events (i.e. absence of evidence about the current state of a stimulus) and explicit absence. Multiple lines of cognitive research are relevant to this question including studies on object permanence, inference by exclusion, associative learning, and information-seeking.

Object permanence can be seen as a basic form of inference about negative information that helps to discriminate between occlusion events and explicit absence. Object permanence was first studied by Piaget (1954) in human children. In the benchmark test of object permanence, the subject witnesses how the experimenter (E) hides the reward under a small opaque box, the displacement device. E then moves the displacement device under one of multiple opaque cups that serve as potential hiding locations. When the displacement device reappears from the visited cup, E shows the subject that the displacement device is now empty. Piaget found that children between 18 and 24 months of age can locate covered objects after so-called invisible displacements. Moreover, developmental research provided evidence for an incremental understanding of occlusion events during the first year of life (based on analyses of infants' looking times; e.g. Aguiar and Baillargeon 1999; Luo and Baillargeon 2005). Following the Piagetian tradition, comparative research examined object permanence abilities in a large number of species (see Jaakkola 2014 for a recent, comprehensive review). However, due to methodological reasons, only great apes (e.g. chimpanzees, Collier-Baker and Suddendorf 2006) and two parrot species (Goffin's cockatoo (*Cacatua goffini*), Auersperg et al. 2014; grey parrot (*Psittacus erithacus*), Pepperberg, Willner, and Gravitz 1997) have so far provided conclusive evidence that they can locate objects after invisible displacements.

A complementary way to study occlusion events (without the requirement to track displaced objects) can be found in the literature on information seeking. Information seeking studies typically address whether subjects seek additional information selectively when they are missing a relevant piece of information (Call and Carpenter 2001). For instance, in the most prominent paradigm, subjects are presented with a number of horizontal, opaque tubes that serve as hiding place for a food reward. In the critical experimental manipulation, subjects either witness the baiting of one of the tubes or not. A number of primate species including great apes, rhesus macaques (*Macaca mulatta*), and brown capuchin monkeys (*Sapajus apella*) seek information before they choose by looking into the opaque containers in particular when they could not see the baiting beforehand (Call 2010; Call and Carpenter 2001; Hampton, Zivin, and Murray 2004; Marsh and MacDonald 2012; Vining and Marsh 2015). Great apes also seek information about the location of a relevant non-food item such as a functional tool (Bohn et al. 2017). Information seeking was also found in the context of a delayed matching-to-sample task in which subjects either got to see the sample stimulus or not. Rhesus macaques and brown capuchin monkeys, unlike pigeons (Roberts et al. 2009), sought information about the sample stimulus when they could not see it previously (Beran and Smith 2011). Whether selective information seeking is supported by meta-cognitive capacities is subject to ongoing debate (Crystal and Foote 2011). Irrespective of this debate, the findings

suggest that there are a number of primate species that compensate absent evidence about the location of objects by seeking additional information.

Another line of research relevant to the question of how animals represent negative information has been conducted with rats (*Rattus norvegicus*). In these studies, rats learned that a stimulus (e.g. a light) predicted a food reward. Following the conditioning procedure, rats were presented with negative information, either with an ambiguous occlusion event (e.g. a covered light bulb) or with an explicitly absent stimulus (a visible but unlit light bulb; Blaisdell et al.2009; Fast and Blaisdell 2011; Fast, Biedermann, and Blaisdell 2016). Rats discriminated between the covered and explicitly absent stimulus and treated the ambiguous covered light bulb more like the present stimulus (the lit light bulb) than the explicitly absent stimulus (unlit light bulb). In a related study, rats distinguished between an occluded reward (a covered drinking receptacle) and an explicitly absent reward in a Pavlovian extinction paradigm (Waldman et al. 2012). Rats again first learned that a light predicted a reward. Then they received an extinction phase in which the light was not paired with the reward any more, either with a covered drinking receptacle (Cover condition) or with an accessible drinking receptacle that did not dispense the sucrose reward (No cover condition). That is, in the No cover condition the reward was explicitly absent whereas in the Cover condition the presence or absence of the reward was ambiguous. In the test after this extinction phase, the cover was removed. Rats' expectations for sucrose when they saw the light were higher in the ambiguous Cover condition compared to the No cover condition, in which the reward was explicitly absent during the extinction phase.

Explanations based on associative learning theory include the so-called renewal effect, which basically states that extinction is context specific (Dwyer and Waldmann 2016). Accordingly, the introduction of the occluder constitutes a change in context which would explain the recovery of a conditioned response toward the occluded option. Given that it remains ambiguous from this hypothesis what constitutes a sufficient change in context, explanations based on the renewal effect are difficult to falsify.

Another associative learning explanation suggests that explicit absence of the stimulus (e.g. the unlit light bulb) acquires associative value (Dwyer and Burgess 2011). When the light bulb is covered this should lead to a general effect on performance (given that the unlit bulb is not visible any more). In contrast, a recent study suggests that when presented with the covered light bulb rats' behavior is governed by retrieved representation of the activated light and not by the fact the rats could not see the unlit light bulb (Fast, Biedermann, and Blaisdell 2016). Non-representational, associative learning based explanations therefore seem to be insufficient to account for rats' performance in these tasks. Discounting these alternative explanations, however, does not mean that rats engage in negation proper when solving these tasks.

Finally, the literature on inference by exclusion suggests that at least great apes, capuchin monkey, and grey parrots are able to infer the location of a reward by discounting alternatives (for reviews see Schloegl, Bugnyar, and Aust 2009; Völter and Call 2017). Typically, subjects are presented with visual or auditory information about the absence of a food reward in a two-choice task. For example, the subject might be presented with a food reward but it cannot witness under which cup the reward is hidden. In the critical condition, subjects get information about the baiting status of the empty cup, either by showing the content of the empty cup to them or by shaking the empty cup (e.g. Call 2004).

While many species can solve the visual version of this task fewer species have been found to solve the auditory version of the task (Völter and Call 2017). In the latter auditory condition, the absence of a rattling sound signals the absence of the reward. Great apes (Call 2004; Hill, Collier-Baker, and Suddendorf 2011), capuchin monkeys (Sabbatini and Visalberghi 2008), grey parrots (Schloegl et al. 2012), and domestic pigs (Nawroth and von Borell 2015) were found to locate the food even when only the empty cup is shaken. Some of the great apes—gorillas (*Gorilla gorilla*) and bonobos (*Pan paniscus*)—and capuchin monkeys even adapted their behavior to the way the silent cup was manipulated by the experimenter: they avoided a silent-shaken cup but not a silent-stirred cup (that should not produce a rattling sound even if it contained a piece of food).

In sum, convergent lines of evidence suggest that at least some nonhuman animals differentiate between different types of negative information. Specifically they discriminate between occlusion events and explicitly absent stimuli. However, all the reviewed examples do not demand negation proper. In most of the test situations, individuals need to consider specific alternatives suggesting that proto-negation might be sufficient to explain their performance. However, proto-negation with present–absent contrary pairs alone might not be a satisfying explanation because it does not explain how individuals discriminate between occluded and explicitly absent stimuli (Blaisdell et al. 2009) or between a silent-shaken and a silent-stirred cup (Call 2004; Sabbatini and Visalberghi 2008). It is hard to falsify such proto-negation explanations because one could easily make up more specific contrary pairs that would explain their performance, such as occluded–absent, or shaken–stirred contrary pairs.

33.2.3.2. *Social reasoning*

Flexible predictions about how others are likely to behave in the future depend, to some degree, on an assessment of the information they have been exposed to in the past (or not). In what follows, we will discuss the relevance of negation for social reasoning in competitive contexts as well as in communication. In order to do so, we assume that, at least some animal species, represent mental states of others. This assumption seems to be warranted in the light of the available empirical evidence (see e.g. Krupenye et al. 2016 as well as Call and Tomasello 2008, Whiten 2013 for reviews), however, some scholars think that mental state ascription is not necessary to explain these findings (Heyes 2015). This discussion is relevant to the discussion of negation in social reasoning as it decides whether one thinks that the types of representations that (proto-) negation operates on are representations of directly perceivable events in the world or the mental states resulting from them. In any case, the question of what type of reasoning process might be used depends, at least according to Bermúdez (2003), not on the content of the representation but on its structure (propositional or not).

A fair number of studies have looked at how animals strategically exploit others' access to certain information in competitive contexts (e.g. Bray, Krupenye, and Hare 2014; Bugnyar, Reber, and Buckner 2016; Dally, Emery, and Clayton 2006; Flombaum and Santos 2005 Schmelz, Call, and Tomasello 2011). For example, great apes have been found to avoid food that a more dominant individual can see or has seen in the past (Hare et al.2000; Hare, Call, and Tomasello 2001). Similarly, great apes avoid looking for food in places that would attract a competitor's attention because accessing them makes noise (Melis, Call, and

Tomasello 2006). Let us describe a fairly recent version of this type of study. Karg and colleagues (2015) conducted a study with chimpanzees in which the experimenter and the participant faced each other, separated by a mesh panel and a small opaque barrier and flanked by two potential hiding locations. These locations were elongated boxes with either an opaque or a transparent lid. Food items hidden in these boxes were accessible from the experimenter's as well as the participant's side. In the beginning of a trial, the experimenter would open the lids towards the participant to let them know which one was transparent and which opaque. Then she placed food items in both boxes and waited. The following interaction was structured so that the experimenter would pull away food items whenever she saw that the participant would reach for them. This was only the case when the participant reached into the box with the transparent lid. Results showed that apes reached into the box with the opaque lid more often than into the box with the transparent lid. On a generous reading, this implies that subjects represent a negated version of their competitor's epistemic state (e.g. 'does not see X') and then choose among the behavioral options available to them which would not lead to consequences that are in conflict with this representation. As mentioned previously, this suggests a wide variety of potential alternatives. In this study, and also the ones cited earlier, individuals made binary choices (e.g. 'reach for left food or right food'). This allows for a, presumably more parsimonious, explanation in that participants represented the two epistemic states held by the experimenter after looking at either the transparent or the opaque lid as a mutually exclusive contrary pair ('sees content of left box or sees content of right box'). By seeing the position of the open lid and inferring that the experimenter would see the content of the box, participants could directly deduce that she would not see the content of the other box (as these two states were mutually exclusive). The subtle consequence of operating on a negated epistemic state of the experimenter (as opposed to a contrary pair) would be the availability of potential alternative strategies in case the primary option becomes unavailable. If, for example, only the box with the transparent lid were available, the same target condition (experimenter does not see content) would also be reached if the experimenter were for example to be distracted or tricked into looking elsewhere. Operating on contrary pairs would make this option not immediately available, as there has been no prior experience that the two states 'experimenter looks at stone I threw to the left' and 'experimenter sees content of the box' are mutually exclusive. While new contrary pairs like this could certainly be formed, they would not be immediately available. We hope that this example illustrates that using negation in social decision making increases the range of available strategies.

Communication differs from competitive social reasoning as described in the last paragraph in that the agent's goal is to actively alter the mental state of another individual by providing some additional information. The relevance of negation for this process lies in the assessment of what the other individual already knows (or not) in order to use communication to bring about the desired change. This pertains more to the speaker side of the interaction and this will therefore be our focus in our discussion below. There are, however, good reasons to believe that communication, at least in humans, depend on hearers' assessment of what the speaker knows (or not) about their mental states (e.g. Tomasello 2008). This additional recursive layer is less important from a negation perspective and so we omit this discussion for now.

On a very basic level, before thinking about how communicative acts alter the mental states of the recipient, speakers need to ensure that hearers can perceive them. Great apes

have been found to adjust their gestures to the attentional state of their recipient (Liebal, Call, and Tomasello 2004; Leavens et al. 2004; Call and Tomasello 2008). They choose more tactile gestures when the recipient is facing away and occasionally move into the recipient's line of sight before producing visual gestures. This echoes the findings discussed above and suggests a basic understanding of the conditions that have to be met for others to have certain mental states, or not. An interesting extension of this is situations in which animals choose to communicate about certain things depending on their recipient's knowledge state. In a recent series of studies, Bohn and colleagues found that great apes (as well as pre-linguistic infants) use pointing gestures to refer to absent entities (Bohn, Call, and Tomasello 2015). In these studies, participants could either point to a visible food item of rather low quality on one plate or point to another, empty, plate that previously contained high-quality food. A point to the empty plate could inform the recipient that one wants another desired food item if they know what that plate previously contained. When Bohn et al. (2016) varied how the participant and the experimenter had interacted around the empty plate before, they found that apes would point more often to the empty plate in cases where there had been some shared interaction around its content before. In the absence of such an interaction, apes almost exclusively pointed to the visible food item. Crockford and colleagues made similar observations when studying wild chimpanzees' alarm calling behavior (Crockford et al. 2012; Crockford, Wittig, and Zuberbühler 2017). When encountering a potential threat, chimpanzees produce alarm calls that inform individuals around them. Such alarm calling is potentially dangerous for the caller as it might attract the predator's attention and callers should only produce them in situations in which recipients are not aware of the threat. In a series of field experiments in which they exposed chimpanzees to dummy snakes, Crockford and colleagues found that individuals would be more likely to produce alarm calls if potential recipients were unaware of the danger because they had not seen it or had been absent during previous calling events. Taken together, these studies suggest that at least some animals tailor their communicative acts to what their recipients are likely to know based on what they experienced. Both examples further suggest that the absence of a certain experience is responsible for the production of communicative acts. One way to construe the corresponding representation is to assume a negated version of the desired state that one wants to change (e.g. 'has not seen snake'). On a proto-negation account, apes might ascribe some kind of positively defined state to their communicative partner which has the same extension (given a certain context) as the negated state described above. In the case of alarm calling this could be something like 'thinks the path is safe' or 'wants to walk along the path'. These states need to be mutually exclusive with the state that the communicative act is intended to bring about (e.g. 'thinks the path is dangerous' or 'wants to walk around the bush with the snake'). Again, the specific contrary pairs that would allow the use of proto-negation in these situations are discrete and formed based on an individual's experience. This implies a very limited set of potential alternatives that could be considered and also makes the provision of fine grained and specific information less likely.

33.2.3.3. *Summary*

In this section we reviewed some of the empirical evidence that relates to the question of whether nonhuman animals engage in negation proper. Many, if not all, of these findings

can be explained in terms of animals reasoning based on mutually exclusive states (proto-negation). These explanations assume that the animal had some specific experiences in the past that led to the formation of each specific contrary pair and/or inferred the mutual exclusivity of certain states based on, for example, some form of causal knowledge. In any case, reasoning based on proto-negation is limited to going back and forth between discrete states of affairs. As we sketched out, the main advantage animals would gain from operating on negation would be the availability of a greater number of alternatives since reasoning based on proto-negation is limited to number of available contrary pairs. The literature on inferential reasoning in nonhuman animals (for a review, see Völter and Call 2017) provides examples in which the performance of animals does not depend crucially on the precise stimulus. Often changing the stimuli or presenting them for the first time does not deter them from making appropriate responses. In some cases, such as the literature on inferences by exclusion, one may invoke non-inferential explanations (e.g. they are avoiding the empty cup) but some new data shows that not even this is a good explanation for their responses (e.g. Jelbert, Taylor, and Gray 2015). Moreover, success in other tasks such as stage 6 invisible displacements or exclusion tasks with two distinct hidden items (e.g. grape and banana) cannot be explained by invoking the avoidance of the empty cup. Instead, subjects' responses appear to be governed by truly inferential processes. In section 33.3 we will build on the idea that negation, in contrast to proto-negation, might allow for more behavioral flexibility and suggest some ways of empirically approaching the question whether nonhuman animals are able to reason based on negated propositional states. We will follow the same structure as above, starting with individual reasoning and then moving on to social reasoning.

33.3. How to distinguish between proto-negation and negation proper empirically

A major challenge for future research is to find methods that would allow us to distinguish between proto- and proper negation. Here we sketch out approaches that could generate evidence that would be suggestive that an individual is negating a proposition as opposed to engaging in proto-negation. None of the approaches below would yield definitive evidence for negation proper, but with sufficient evidence accumulating, negation proper might turn out to be the better explanation for the data at hand. Most of the suggestions below refer to the idea that operating on negated propositional states (as opposed to mutual exclusive states in the world) increases the number of available behavioral responses. Furthermore, given that negation in humans is a logical operation, it is independent of the content of the propositions that are being negated. Therefore, even if this kind of evidence were found in some animal species, it would be necessary to show that the ability is not restricted to a few specific situations (e.g. reasoning about tool functionality or object permanence).

33.3.1. Individual reasoning

One potential way to distinguish between proto-negation and negation proper would be to look at individual differences (see Völter et al. 2018 for a recent plea for looking at individual differences in comparative cognitive research). As mentioned above, the respective contrary pairs that figure in proto-negation are experience dependent and therefore specific to certain situations. This implies that individuals differ in their proto-negation abilities due to the kind and amount of experience they possess, and not due to variation in a domain-general cognitive ability. Based on the proto-negation account, individuals who do well in one task should not necessarily do well in another task. On the other hand, in the case of negation proper, one applies the same logical operation to all sorts of propositions. It may therefore be regarded as a domain-general cognitive ability. This implies systematic individual differences across tasks. When testing individuals in a series of tasks involving reasoning from one state of affairs to its negated counterpart(s), the two accounts make differential predictions. Finding that animals vary systematically in their performance would at least be suggestive evidence that the same kind of logical process is involved in all cases. Importantly, consistency in performance in such tasks could originate from systematic differences between individuals in learning/inferential abilities that lead to the formation of contrary pairs. For example, Herrmann and Call (2012) found that some great apes consistently perform better than others in a variety of cognitive tasks. These individuals might be better at spontaneously inferring mutually exclusive states and would therefore also perform better in tasks requiring reasoning from one state of affairs to its negated counterpart. Nevertheless, finding no relation between tasks would speak against a common logical basis.

33.3.2. Social reasoning

When discussing the literature on competitive social reasoning, we mentioned that most studies to date asked subjects to choose between one of two responses. Making the correct decision in these scenarios is well captured by the idea that animals reasoned based on contrary pairs. This conclusion could be challenged if participants were found to manipulate the experimenter's epistemic states in multiple different ways and chose flexibly among alternatives. For example, we could imagine a scenario in which an experimenter is guarding a reward and whenever she sees the subject approaching it, she takes it away. To obtain the reward, the subject would either have to actively distract the experimenter or approach the reward in such a way that the experimenter cannot see it (see e.g. Whiten and Byrne 1988, Coussi-Korbel 1994, Hirata and Matsuzawa 2001 for different examples of such tactics). By varying features of the setup, such as the presence of barriers or potentially distracting objects in the vicinity, from trial to trial, one could see if the subject adjusted flexibly to each situation. For example, crouching behind a barrier when available or throwing a ball behind the experimenter when there is one. The common thread connecting these alternatives would be that the animal is trying to keep the experimenter from entertaining a certain epistemic state (e.g.: 'sees my hand when I reach for the food') and alternatives might be chosen because they do not have this state as an anticipated

consequence. Flexible switching between alternatives would at least be suggestive that the desired effect of each action is negatively defined and therefore represented as a negated proposition instead of a positively defined state of affairs in the world. Interestingly, there is some evidence that apes flexibly switch between alternatives in the case of a positively defined state of affairs (i.e. 'X looks at me', see e.g. Liebal, Call, and Tomasello 2004; Leavens, Hostetter, Wesley, and Hopkins 2004).

A structurally similar approach could be taken in the area of communication. In the studies mentioned above (Bohn et al. 2016, Crockford et al. 2012), apes decided between producing a given communicative act (pointing to the empty plate, making an alarm call) or not. In line with a proto-negation explanation, the mental state that they ascribed to their communicative partner and which guided their decision whether or not to communicate could have been defined in a positive way as being about some concrete state of affairs in the world. From experience, they knew that a given communicative act would lead their partner to entertain the complementary (mutually exclusive) state they desired them to have. The ascription of a negatively defined, propositional state could be the more parsimonious explanation if there were evidence that animals flexibly provide information that their communicative partner lacks in situations that require multiple, complementary bits. For example, if information A and B both need to be known, they would provide A to someone who does not know A and B to someone who does not know B. It would still be possible that animals have experience with and therefore represent each of the four combinations of information A and B (AB, A¬B,¬AB, ¬A¬B) as discretely different states, it would just be less plausible, especially when more bits are added. We are aware that the approaches outlined here require additional cognitive capacities that might be beyond what animals are capable of. For example, spontaneous and flexible switching between responses requires subjects to construct these alternatives and simulate their likely effects on the experimenter. This requires a fairly elaborate theory of the experimenter's mind. Nevertheless, we hope that these examples serve to illustrate possible ways to study how animals cope with situations that would typically involve negation in humans.

33.4. Negation in the evolution of human reasoning

We argued that the main benefit of reasoning based on negation, as opposed to proto-negation, is the larger number of conclusions that can be drawn from a given premise. Proto-negation is limited to discrete contrary pairs whereas the conclusions to be drawn from negated statements are potentially limitless (in the way that the opposite of white is not black but non-white). The main prerequisite necessary for negation proper is propositionally structured thoughts. From an evolutionary perspective, emergence of negation had to be preceded, or at least accompanied, by the emergence of propositional thought (Bermúdez 2003). It is beyond this chapter to give a comprehensive overview of all the accounts on how propositional thought evolved. One way to think about this would be a Vygotskian view (2012, see e.g. Moll and Tomasello 2007) in which propositionality emerges first in communicative interactions as a consequence of multiple agents

coordinating their, sometimes differing, perspectives on one particular state of affairs. Differing perspectives encourage communicative partners to think about the specific way in which someone else sees something and creates a need to represent the corresponding intension. Such an account presupposes the emergence of fairly sophisticated social-communicative skills. However, once ways to represent the intensionality attached to a particular perspective in communication emerged (in vocalizations or gestures), this form of representation could subsequently be internalized and used for thought. Negation might emerge in conjunction with propositional thought because individuals need to be able to reject each other's particular perspectives in communicative interactions. Such a rejection expresses that one thinks that a particular perspective is false as opposed to a particular state of affairs not being the case. For example, if you tell me (e.g. by pointing and pantomiming) that there is a bear hiding behind the bush but I saw the bear walk off earlier, when I reject your statement, I do not reject that there is a bush but your perspective that there is a bear behind the bush. The emergence of negation and propositional thought would be intimately linked in this case.

33.5. Conclusion

To date, there is no compelling evidence for negation proper in nonhuman animals. This conclusion is based on a particular perspective on negation, namely that it involves the negation of a particular proposition and is therefore tied to propositional thought more generally. As we have seen, there is plenty of evidence suggesting that animals engage in forms of reasoning that are functionally related to negation but might be procedurally different. Nonhuman animals might engage merely in proto-negation (Bermúdez 2003), that is reasoning based on mutually exclusive contrary pairs. However, behavioral evidence is often ambiguous with regard to the underlying cognitive processes and there is accumulating evidence for high degrees of behavioral flexibility in animals also in situations involving negative information (e.g. about the absence of a stimulus) and so the jury is still out. Joint efforts from comparative psychologists, philosophers, and computational cognitive scientists will allow us to specify how and when negation (and propositional thought) emerged in phylogeny.