University Press Scholarship Online
Oxford Scholarship Online

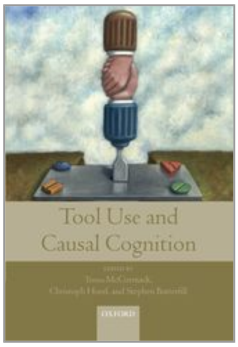

Tool Use and Causal Cognition
Teresa McCormack, Christoph Hoerl, and Stephen Butterfill

Print publication date: 2011
Print ISBN-13: 9780199571154
Published to Oxford Scholarship Online: January 2012
DOI: 10.1093/acprof:oso/9780199571154.001.0001

Causal Knowledge in Corvids, Primates, and Children

More Than Meets the Eye?

Amanda Seed
Daniel Hanus
Josep Call

DOI:10.1093/acprof:oso/9780199571154.003.0005

Abstract and Keywords

Previous experimental work has led several authors to conclude that only humans reinterpret ‘first-order perceptual relations in terms of higher order role-governed relational structures,’ meaning that among other things, non-human animals are not capable of physical reasoning based on abstract, unobservable object properties. Instead they must rely on first-order perceptual information to solve problems. Such an account could approximate behaviour driven by physical knowledge very well if natural selection has pre-prepared the learning animal to attend to perceptual features of the stimulus that are most likely to correlate with its functional properties in the natural environment. This chapter aims to challenge the ‘relational reinterpretation hypothesis’ put forward by Penn et al. (2008) by describing data from several other paradigms aimed at distinguishing between explanations based on surface-level perceptual characteristics and those in which object properties (such as solidity, continuity, weight, and rigidity) are represented at a deeper level of abstraction, where ‘abstract’ means that the information is not equivalent or reducible to concrete, analogue sensory input, but rather has undergone further processing in which meaning is extracted.

Keywords: first-order perceptual relations, physical reasoning, abstract, object properties, relational reinterpretation hypothesis, perceptual information

The majority of mobile animals need to locate their food in space (and in some cases, also time). In addition they may need to assess the quality and quantity of located resources to make efficient foraging decisions, and some food may need to be extracted or processed before it can be consumed. The environment can therefore pose three kinds of challenge to foraging animals: to locate, quantify, and extract food (Tomasello & Call, 1997). The ability to learn allows animals to face these challenges—especially in changing environments—and exploit variable sources of information that a hard-wired mechanism would struggle to process. The argument has been made from psychologists such as Thorndike (1898) that such learning can be based on covariation without any need to invoke the concept of causation (e.g., animals could learn to search for food under fruit trees when the wind blows, simply by using the temporal contiguity between gusts of wind and fruit availability to associate the two events). However, an ability to extract and encode the causality underpinning the surface information received by the senses could facilitate the difficult task of sorting out the relevant information from the many events that covary—leading to appropriate and flexible future behavior. Seed and Call (2009) made a distinction between two aspects of causality. First, events are underpinned by predictable *causal structures* (e.g., gusts of wind cause branches to shake, and not the other way round), and second, objects have *causal properties*: they obey physical laws that constrain the possible ways in which they can interact with one another (e.g., solid objects cannot pass through one another, and fall if they are unsupported).

Cognition of both aspects of causality has been investigated in non-human animals and developing children. Gopnik et al. (2004) have argued that children develop a causal map of the events in the environment: “an abstract, coherent, learned representation of the causal relations among events.” There is some evidence that non-human animals also go beyond association and encode causal directionality, or “*what* caused what” (Blaisdell et al., 2006; Dickinson & Shanks, 1995; Waldmann et al., 2006), especially in the case of their own actions’ causal relationship to outcomes. On the other hand there is the ability to cognize higher order causally relevant object properties, the “*how* it caused it.” This second aspect will be the focus of this chapter, in which we will review evidence for object knowledge and physical reasoning in corvids, non-human primates, and developing **(p.90)** children from two sorts of tasks: reward extraction and reward location. Each taps into a challenge posed by the natural environment, and so we hope is an ethologically valid way for assessing what subjects know about causality. The first asks subjects to extract a reward that is out-of-reach. Traditionally this type of paradigm has employed tool-using tasks, which hinge on the relation among the following elements: the tool, the reward, and some obstacle. Often the task consists of anticipating the effects that certain properties of the elements involved will have on the final outcome. The second paradigm asks subjects to locate a reward. Compared to classical tool-using tasks, they represent a more recent addition to the methods devoted to investigating object knowledge. One key feature of this group of tasks is that in order to infer the location of the reward subjects must take into account certain object properties. Although the extraction and location tasks differ substantially in their implementation and measures—one is based on producing object interactions while the other is based on observing object interactions—they can play a complementary role (even on the same task) as will be illustrated in what follows.

As we will describe, previous experimental work has led several authors (most prominently, Povinelli and colleagues) to conclude that only humans reinterpret “first-order perceptual relations in terms of higher order role-governed relational structures,” meaning that among other things, non-human animals are not capable of physical reasoning based on abstract, unobservable object properties. Instead they must rely on first-order perceptual information to solve problems. Such an account could approximate behavior driven by physical knowledge very well if natural selection has pre-prepared the learning animal to attend to perceptual features of the stimulus that are most likely to correlate with its functional properties in the natural environment. In this chapter we aim to challenge the “relational reinterpretation hypothesis” put forward by Penn et al. (2008) by describing data from several other paradigms aimed at distinguishing between explanations based on surface-level perceptual characteristics and those in which object properties (such as solidity, continuity, weight, and rigidity) are represented at a deeper level of abstraction, where more “abstract” means that the information is not equivalent or reducible to concrete, analogue sensory input, but rather has undergone further processing in which meaning is extracted.

We suggest that the results of these experiments require an alternative explanatory framework to the distinction made by Penn et al. (2008) between first-order perceptual knowledge and abstract relational knowledge, particularly because this view confounds, within the perceptually based category, the different ways in which perceptually detectable features may be represented, for example, reasoning based on physical laws, such as that a solid object cannot pass through a barrier, and reasoning based on the spatial relationship between a reward and a feature of a certain shape or color. Within the abstract category, knowledge of physical properties is conflated with purely symbolic abstractions such as words and concepts that are not grounded in natural categories, and whose acquisition depends on indisputably unique forms of language, culture, and teaching. To move beyond the conceptual dichotomy suggested by Povinelli, Penn and colleagues we favor a framework based on three progressively deeper levels of abstraction for representing the causal information in the environment: *perceptual*, *structural*, and *symbolic*. This is because of the evidence that, in some contexts, animals such as corvids and apes use abstract structural knowledge of object properties within a causal framework to (p.91) solve problems. Indeed, we show that for apes, learning based on arbitrary perceptual cues is difficult. In support of the further distinction between abstract physical knowledge and symbolic reasoning, we show that in contrast to the use of structural knowledge in early infancy, the ability to use arbitrary cues or symbols to solve problems emerges fairly late in child development.

1. Extraction problems

Traditionally, much of the work on causal knowledge in primates has focused on tool use. Researchers have used relatively simple tasks such as the support or the stick problem in which subjects must retrieve an out-of-reach reward by using a tool (Natale, 1989; Piaget, 1952; Spinozzi & Potí, 1989). More complex tasks require subjects to overcome some obstacle (not just the distance) between the reward and the tool. One of the most well-known tasks of this kind is the trap tube (Limongelli et al., 1995; Visalberghi & Limongelli, 1994). In this section we will review data on the support/connectedness task and variations on the trap tube task.

Support/connectedness task

When presented with the connectedness task—two objects attached to food rewards (strings, strips of cloth or pre-positioned tools), one intact and one with a clear break in the middle—great apes, vervet monkeys, cotton-top tamarins, elephants, parrots, and pigeons are able to pull the connected, continuous object to bring the food within reach. However, whilst some species performed significantly above chance from the start of the experiment in at least some configurations (great apes, vervet monkeys, and elephants), pigeons and cotton-top tamarins required extensive training (over a hundred trials before the correct solution was learned). It is tempting to infer that the species that solved the problem spontaneously used a qualitatively different cognitive mechanism, involving an appreciation of the principle of connectedness such as would underpin an adult human's behavior, whilst the other species relied on a simpler mechanism based on learning an arbitrary association between the appearance of the correct alternative and reward. However, an animal's performance depends on a number of both cognitive and non-cognitive processes. For example, those animals that take longer to solve a particular task, or even fail it completely, may be less motorically dexterous, less motivated, more easily distracted, find the task at hand harder to perceive, or find irrelevant features of the task more attention-grabbing, compared to the species that solve it quickly (Bitterman, 1975). Even the same individuals can perform very differently on two tests supposedly probing the same ability. For example, although the great apes tested by Herrmann et al. (2008) on the connectedness problem were able to solve it spontaneously when the material involved was string or cloth, they performed at chance (in the six trials given) when the objects were two wooden canes, pre-positioned around the food rewards. Had only one configuration been given, a very misleading picture of the animal's abilities would have emerged. Of course, even an identical setup may not be equivalent for different species. Given that many factors can impact on a subject's performance, it is clear that caution must be exercised when interpreting negative results and species differences.

Of course even when subjects solve the task immediately, they may not have encoded anything about object properties, but simply relied on the tasks' surface appearance, either because they preferred the correct alternative due to inborn predispositions or biases or because pulling objects with a similar appearance was previously associated with reward. To investigate this possibility, Hauser and colleagues gave transfer tests to cotton-top tamarins that varied functionally irrelevant features such as the cloth's color, shape, texture, and the shape and size of the gap (Hauser et al., 1999). The tamarins readily transferred their solution across the majority of these changes, suggesting that they had used functionally relevant properties to solve the original discrimination. Importantly, pigeons that had learned to solve one version of the connectedness task failed to transfer to a new version in which the shape and color of the material was changed, suggesting that they had relied on functionally irrelevant perceptual cues to solve the original task (Schmidt & Cook, 2006). Nevertheless, in the absence of verbal report, even when an animal transfers its solution to a new context it is possible that its behavior can still be explained by generalization based on surface-level perceptual characteristics, especially if it has evolved biases to attend to the particular perceptual dimensions of a feature that coincide with its functionally relevant aspects in the contexts that the animal usually experiences it. Indeed, infant tamarins with little experience of using tools attend to functionally relevant features such as size and shape rather than irrelevant changes in texture or color when selecting a tool (Hauser et al., 2002). Almost all of the transfer tests in the connectedness task can be solved using a common perceptual metric, namely, avoid the gap between yourself and the reward. In experiments conducted with nursery reared chimpanzees, Povinelli (2000) made the differences between the two alternatives more subtle, and his subjects failed some telling transfers. For example, when one of the tools was broken but the ends were aligned in front of the chimpanzees, they did not avoid the broken tool, even though the properties of the two options had been demonstrated to them. It seems that chimpanzees may have difficulty in making the discrimination when perceptual information concerning the lack of continuity is not available at the time of choice. Does this mean that they did not use a notion of connectedness to solve the original condition? Or could limitations in other faculties (attention to the demonstration, working memory, inhibition) have played a role in their poor performance over these few trials?

We recently took a different approach to the question. It may be the case that discriminative perceptual cues are necessary for successful performance, although this question warrants further study. But are they sufficient? To address this question Seed et al. (in preparation) compared performance on two versions of the connectedness task. In the standard “transparent” version the rewards were tied to strings, one complete string and a broken one with a 5-cm gap in the middle. In the “covered” version the table on which the strings rested was covered by a lid, such that though the rewards could be seen at the distal end of the table, and the ends of the strings could still be seen and pulled at the subject's end, the central part of the real strings was obscured. We stuck a broken and an unbroken string to the cover in the same place as their real counterparts beneath, so that subjects could still use the appearance of the break in the middle of a piece of string to correctly avoid pulling from that side, or trace an unbroken white line (p.93) to the food to choose the correct alternative (Figure 5.1). We tested six chimpanzees and six bonobos; half received the transparent task first, half received the covered version. Five of the chimpanzees and three of the bonobos solved the transparent version of the task, some from the first or second day of testing. However, none of the subjects solved the covered version in 100 trials. There was no effect of the order in which they received the tasks, nor was there an effect of the location of the break in the string (near to the ape or near to the food). This means that even subjects that were using the gap in the real strings to correctly select the unbroken one did not use the gap shown on the cover, even though perceptually this was a very similar cue.

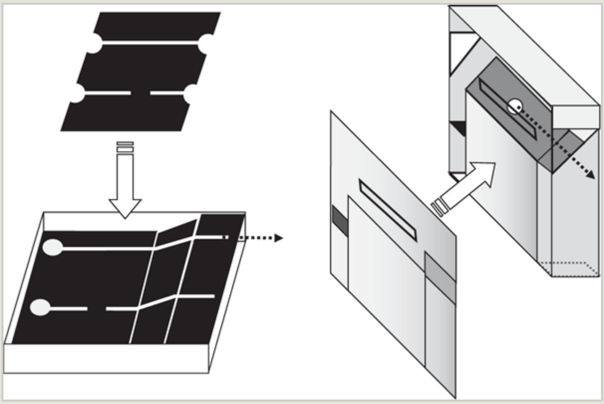

Figure 5.1 Connectedness task (left) with cover above and trap box task (right) with cover to the side. The dashed arrows indicate the correct solution.

This study aimed to test the null hypothesis that successful performance on the broken string task reflects rapid learning based on the perceptual difference between the two alternatives and differential reinforcement. We argue that this cannot account for the performance of chimpanzees and bonobos on the broken string problem. Otherwise, they should have learned under both conditions since both displays were very similar and pulling the correct string produced the same feedback. Preschoolers performed similarly: seven out of twelve 3½-year-olds and all of the older children we tested solved the transparent task, but performance on the covered condition was very different. None of the 2½-, 3½-, and only one of the 5½-year-old children solved the task if they were presented with the covered condition first. Two out of six 3½-year-olds and five out of six 5½ and 6½-year-olds solved the covered version if they had already solved the uncovered task, but when they received the covered condition first, only three out of six 6½-year-olds could use the cues stuck to the lid to solve the task. It seems that children, like chimpanzees and bonobos, use knowledge of object properties such as continuity or connectedness to solve problems. The ability to use arbitrary cues with no obvious causal relevance to the task seems to emerge much later in development. This may reflect the fact that the covered condition requires the interpretation of a symbolic cue, a point to which we will return later.

(p.94) Trap task

Another task that has been used extensively in comparative studies of animal causal reasoning is the trap task. In the original version of the task, subjects are faced with a transparent tube that has a trap in its centre and a reward placed out-of-reach of the subject inside the tube next to the trap. Subjects are given a stick whose diameter is slightly smaller than the inner diameter of the tube. To solve the problem subjects should insert the stick inside the tube and push the reward away from the trap. This task has proven extremely difficult to solve for capuchin monkeys and chimpanzees as only a minority of subjects solved this task even after dozens of trials (Limongelli et al., 1995; Povinelli, 2000; Visalberghi & Limongelli, 1994). Furthermore, when the tube was inverted so that the trap was no longer functional, subjects continued to avoid it, suggesting that they had not understood anything about causal properties, but had rather formed a rule based on the perceptual features of the trap.

Some variations on the trap tube aimed at simplifying the task have produced similar results (see Call, 2010, for a review; see Figure 5.2). For instance, Povinelli (2000) presented chimpanzees with a pair of rakes each with a reward in front of them. Crucially, one rake also had a trap in front of it while the other simply had a painted patch of the same dimensions of the trap (i.e., fake trap). In order to succeed, the only thing that subjects had to do was to pull the rake placed behind the fake trap since pulling the other rake invariably sent the reward into the trap. All subjects except one failed to solve the task, and even this subject did not pass subsequent control tasks—something that has also been observed in the original trap tube task (Limongelli et al., 1995; Povinelli, 2000). This has led several authors to conclude that subjects may have used a perceptual strategy based on using the position of the trap to determine the appropriate insertion point but without understanding that the position of the reward with respect to the trap hole is the critical feature (p.95) in this task. Povinelli (2000) concluded that apes had a limited understanding of the physical properties of the trap. Subsequent studies carried out with capuchins and gibbons used a similar paradigm, and also concluded that subjects did not have a total comprehension of the elements of the problem but that subjects might have learned certain associative rules (Cunningham et al., 2006; Fujita et al., 2003).

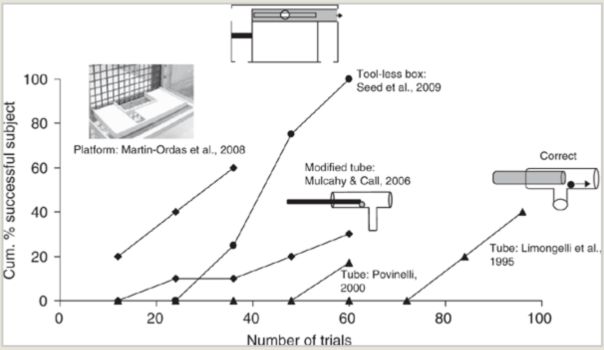

Figure 5.2 Trap tube setups and the cumulative percentage of subjects who performed above chance as a function of the number of trials that they received in each study.

However, the results improve dramatically after implementing certain modifications such as allowing subjects to choose what actions to use (pulling or pushing) (Martin-Ordas et al., 2008; Mulcahy & Call, 2006), where to insert the tool (Girndt et al., 2008), or if the need to use a tool is completely eliminated (Seed et al., 2009). For instance, Girndt et al. (2008) tested twenty chimpanzees on the trap table problem, and reported that apes performed significantly better when they were given one tool to use, rather than a choice between two pre-positioned tools (Figure 5.3). When tested using Povinelli's (2000) original setup, the chimpanzees performed at chance in the first ten trials, just as most of Povinelli's chimpanzees did, but when they were tested with just one tool, 80% of subjects raked from the correct side in their first trial. Mulcahy and Call (2006) made the trap tube wide enough for the chimpanzees to push or pull the reward with the tool. Nine out of ten apes preferred to pull. Not only did the three successful subjects (two orangutans and one chimpanzee) learn much faster in this study than in previous ones (they took an average of forty-four trials to reach criterion), but all of these subjects passed the inverted control task. Tebbich and Bshary (2004) found that a woodpecker finch, a tool-using species of bird from the Galapagos, could also solve this version of the task. Tebbich et al. (2007) tested rooks, a species of corvid that does not usually use tools in the wild, on a version with a pre-inserted stick (pulling the stick (p.96) would move the food) and found that three out of eight subjects could solve it. Seed et al. (2009) cut holes in the front of the problem wide enough for subjects' fingers but not wide enough for the reward, and found that all of the eight chimpanzees tested solved the task, taking an average of 66 trials to reach criterion. These results are important because they highlight that some of the difficulties experienced by the subjects in solving the task may

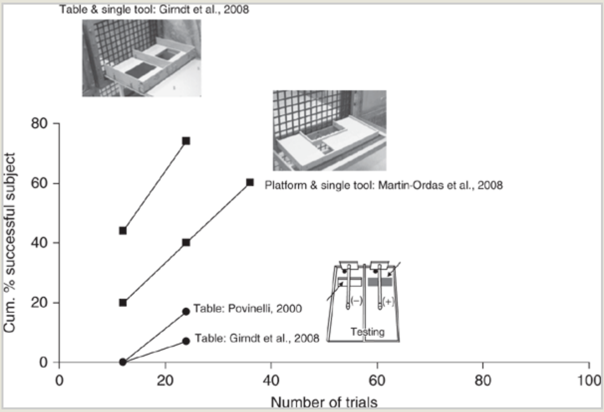

Figure 5.3 Trap table setups and the cumulative percentage of subjects who performed above chance as a function of the number of trials that they received in each study.

have more to do with the particular task implementation than a lack of knowledge about the functional properties of traps. They also show that success on this task is not limited to primates, or to tool users.

The second problem with the traditional setup is interpreting results from the inverted control task. Negative results, when subjects continue to avoid the trap, are inconclusive because there is no cost to continuing to use this strategy. Positive results, when subjects revert to random responding, are also problematic, because to argue for a representation of causal properties from this evidence would mean a strong interpretation of a null effect. Furthermore, subjects could just have been using a very specific perceptual cue in the initial test phase. Seed et al. (2006) aimed to address these problems. The aim was to test the null hypothesis: “a successful animal will use an arbitrary cue to solve the task.” Eight rooks were tested on a version of the trap problem that featured two “traps” along a horizontal tube. One of the traps was functional (sealed with a black disc at the bottom) and would trap the reward if the rooks pulled the food over it. The other was non-functional; in Design A it had a black disc at the top, which the food could pass over; in Design B it had no black disc, so the food could fall through it. Seven of the eight birds learned to avoid the functional trap, in between 30 and 140 trials.

All seven rooks immediately solved task B once they had learned to solve A, and vice versa. However, both of these tasks could have been solved by learning to avoid the trap with the black disc at the bottom, without anything about the properties of the objects being encoded. The seven birds were therefore given two transfer tasks, both featuring the two previously non-functional traps (pass-across or fall-through). In Design C both ends of the tube were blocked with bungs, so the food could not be recovered from the end of the tube, and the birds needed to pull away from the trap with the black disc at the top; in Design D the tube was lowered to the surface of the testing shelf, so that the food could not be recovered from beneath, and the rooks needed to pull toward the trap with the black disc at the top to be successful. Crucially, therefore, both tasks featured the same familiar cue, but each required the opposite response to it (pull away from the black disc in

Task C, pull toward it in Task D). The birds were given 20 trials on both of these transfer tasks. Six of the subjects performed at chance on both tasks, but one bird was able to solve these transfers, suggesting that she did not solve the two-trap task simply by using the appearance of the functional trap as an arbitrary, surface-level cue (Seed et al., 2006).

Seed et al. (2009) recently conducted a similar experiment with chimpanzees, in which subjects could move the food reward with their fingers rather than a tool. The setup was slightly different, because the task was a box with a transparent Perspex front (see Figures 5.1 and 5.2). Nevertheless, four designs were used, equivalent to A–D in Seed et al. (2006). Like the rooks, the chimpanzees were first given 100 trials to learn to solve designs A and B, before being given designs C and D as transfer tasks. One chimpanzee passed both designs C and D. Like the successful rook, this chimpanzee could not have been using a rule based on an arbitrary perceptual cue to solve the task.

(p.97) We compared the performance of these experienced chimpanzees to naïve ones on a new version of the task, which differed from the original task in size, shape, color, and material (it resembled the two-trap-tube test used to test rooks). Strikingly, the experienced subjects solved the task rapidly, but all but one of the inexperienced subjects failed to do so in 150 trials. Similarly, Taylor and colleagues (2009) recently found that three out of six New Caledonian crows also learned to solve the two-trap problem, and that successful subjects were able to transfer to the trap table task (similar to that in Figure 5.3), which was as different from the original problem as the second task given to chimpanzees was. These results suggest that the rook, chimpanzees, and New Caledonian crows did not use simple perceptual cues to solve the trap task. We propose that instead they extracted causally relevant functional information (such as surface continuity, or the solidity of barriers). However, further work is required to uncover the exact nature of their object representations, and the algorithms by which they are fed into behavior. The individual differences in all of these experiments are striking, and also warrant further attention.

Seed and Call (in preparation) recently tested children on the trap box problem described earlier. In this study, we used an approach similar to that described for the connectedness task to see whether perceptual cues provide sufficient information in the absence of their causal relevance: we tested one group of children on the transparent version (similar to that given to chimpanzees) and another group on a covered version, in which stickers were placed in the same location as the traps and barriers to provide reliable cues (Figure 5.1). In the transparent version, four of eight 2½-year-olds, and all eight 3½-year-olds passed the trap task, and interestingly, all but one of the successful children passed the transfer tests C and D. However, none of the children tested at this age solved the covered version in the 40 trials given. Similar to the results from the connectedness task, it was only much older children that could use the arbitrary cue to solve the task: none of the eight 4½-year-olds, three of the eight 5½-year-olds, and eight of the sixteen 6½-year-olds were successful. Interestingly, two of the unsuccessful subjects performed significantly worse than expected by chance, which means that they responded to the cue, but they almost always moved the reward toward the sticker in the trap position (and therefore lost the reward). Some subjects that solved the initial task went on to fail the transfer tasks, something that never occurred on the transparent task (on which we also tested 6½-year-olds). The inability of young children to solve both this task and the covered connectedness task (Figure 5.1) may reflect the fact that the task rests on the interpretation of a symbolic (iconic) cue. Several studies have documented that becoming “symbol minded” is a significant hurdle for young children during development, because using a symbol requires recognizing its dual nature (DeLoache, 2004). For example, children do not use the hiding location of a small toy in a scale model of a room to locate the toy in the real room until over 3 years of age, and in the absence of explicit verbal instruction about the model room relation even 5- to 7-year-olds struggle. In this task, the blue cue is a piece of plastic that cannot possibly impede the passage of a ball. However, it is also a symbol, which tells the child which way to move the ball in order to get it out of the box, but to recognize this may mean suppressing intuitions about its causal irrelevance based on its physical properties. This study suggests that, for young children, knowledge of object properties may be essential for solving problems (p.98) quickly and flexibly. Using arbitrary cues with no obvious causal relevance to the task at hand seems to be a much steeper cognitive challenge, and one that emerges much later in development, especially in the absence of verbal instruction.

2. Location problems

The inferential abilities of primates have been documented extensively in various domains (see Call & Tomasello, 2005; Tomasello & Call, 1997 for reviews). Monkeys and apes can infer the location of hidden objects based on either arbitrary associations between stimuli (e.g., transitive inference: Boysen et al., 1993; Gillan, 1981) or spatiotemporal constraints in the test situation (e.g., object permanence: Call, 2001; de Blois et al., 1998; object individuation: Mendes et al., 2008; Phillips & Santos, 2007; Santos et al., 2002). Much less is known about inferences based on object–object interactions, other than the fact that chimpanzees and rhesus macaques associate certain object transformations with particular outcomes (e.g., cut apple with knife, Hauser & Spaulding, 2006; Premack, 1983). Investigating the use of object–object interactions for inferential purposes is particularly interesting because animals experience a variety of object–object relations in their everyday lives. From the point of view of this discussion, we want to know whether animals can take into account the effect that objects will have on one another in order to infer the location of food, and more importantly, whether they use information about the physical properties of the objects involved to do so.

Support

Let us begin with an example that links with the trap task section. Martin-Ordas and Call (2009) presented apes with a platform that had two square holes cut on it so that it created three solid areas on the front part of the platform: one central area and two smaller areas next to each hole on each side of the platform (see Figure 5.4). One hole was covered with a transparent piece of plastic and the other hole was left uncovered. Two opaque plastic cups are placed upside down side by side on the central area of the platform next to the holes. The experimenter showed a reward to the subject and behind a screen placed it under one of the cups so that the subject did not see its final destination. After the baiting was completed, the experimenter removed the screen and laterally displaced each cup from the central area to the side so that each cup crossed over the hole closest to them. After both cup displacements were completed, the ape could select one of the cups by touching it. In order to avoid the noise that the reward would make when it fell through the open hole, we never displaced the reward over the open hole but it was always displaced over the covered hole.

We found that apes selected the baited cup above chance levels both overall and in the first trial but failed to do so if both holes were covered with opaque or transparent pieces of plastic (Figure 5.4). This ruled out the possibility that subjects used inadvertent cues left by the reward or the experimenter to solve the problem. Apes also failed to select the baited cup if the displacements occurred when both holes were covered but later one hole was covered with a transparent piece of plastic and the other was left uncovered. Since this is the same perceptual configuration that subjects encountered at the time of choice under the experimental condition, we can rule out that subjects had (p.99) a predisposition for avoiding uncovered holes regardless of the reward displacements. It also indicates that subjects were not choosing based on the final configuration alone (i.e., simply avoiding cups next to the hole). Moreover, Martin-Ordas and Call (2009) found a positive correlation in apes' performance between the gap task and some trap tasks reviewed in the previous section.

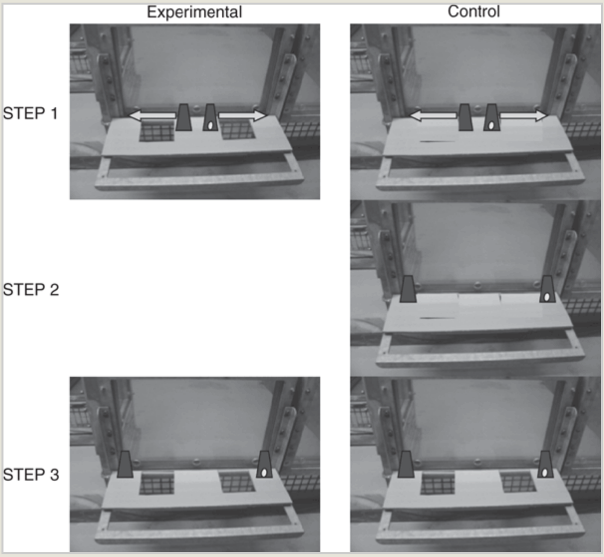

Figure 5.4 Experimental setup and procedure for the gap task (Martin-Ordas & Call, 2009). In the experimental condition, one of the platform gaps is covered by a transparent piece of Plexiglas while the other is left uncovered. Behind a screen a reward is hidden under one of the two centrally located cups. Upon removal of the screen, each of the cups is successively dragged to the left or the right until they reach the end points on opposite sides of the platform. The baited is always dragged over the covered hole; otherwise, the food would fall. After both displacements are completed, the subject can pick one of the cups. In the control, the procedure is the same except that gaps are covered by opaque Plexiglas pieces and after the displacements have been completed one opaque Plexiglas piece is replaced with a transparent one, whereas the other gap is left uncovered so that before subjects choose, the setup is identical to that of the experimental condition.

Weight

Although the results of this experiment were clear, one may wonder whether this pattern of results is peculiar to this setup. In other words, will apes respond differently to causally relevant cues and arbitrary cues in tasks involving causal principles other than support? Let's examine two other studies aimed at the same question but with different arrangements. These experiments investigate specifically whether subjects can infer the location of food (p.100) based on its weight, or better the effect that its weight has either on other objects or on their own body.

Chimpanzees implicitly take into account the effect that their own weight has on pliable vegetation when planning their climbing actions (otherwise, they would fall from trees all the time). However, from the mere observation of such behavioral adaptations it remains unclear whether it is the outcome of rather inflexible heuristics (e.g. always avoid branches below a certain dimension) or whether chimpanzees are truly sensitive to the causal interrelation between their own weight and objects in the environment. Furthermore, very little is known about whether they can use the weight of external objects to make inferences about the location of hidden rewards. Hanus and Call (2008) presented chimpanzees with two opaque cups mounted on opposite sides of a balancing beam kept in equilibrium by a pivot located under its centre of gravity. In one condition the experimenter hid a reward inside one of the cups and released the beam, which resulted in the baited cup moving downward and the empty cup moving upward. Once the beam had reached this new equilibrium, subjects were allowed to select one of the cups. Obviously, the correct cup was always the lower cup. We compared the causal condition with two control conditions. In the static control condition we assessed whether subjects preferred to select the lower cup rather than the upper cup when they were mounted on a static inclined beam so that the weight of the reward was not responsible for the fixed orientation of the cups. In the external cause condition, the setup was identical to the causal condition except that when the experimenter released the beam after baiting, it maintained its horizontal equilibrium until the experimenter pushed physically down the beam to its final slanted orientation. Thus, the experimenter, not the weight of the reward, was responsible for the change in orientation. Chimpanzees selected the baited cup in the causal condition but not in any of the control conditions (see Figure 5.5). Moreover, subjects that were performing at above chance levels under (p.101) the experimental condition responded at chance levels upon receiving the static control condition. Conversely, subjects that were responding at chance levels under the static control condition began responding above chance under the causal condition. Equally remarkable is the difference between the causal and the external cause condition given that the information about the beam's displacement and the final position of the baited cup were identical under both conditions. These data also help us rule out the possibility that subjects have a predisposition to respond to certain stimuli in certain ways. For instance, the external cause control has the same movement and start- and end-positions as the causal balance condition. Similarly, several of the control conditions in the trap task showed that subjects did not have a general predisposition to avoid holes or traps.

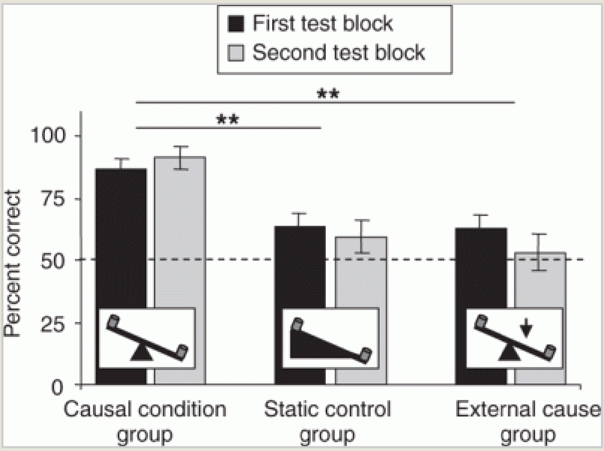

Figure 5.5 Percentage of correct trials under each of the conditions of the balance beam experiment by Hanus and Call (2008). Reprinted from Current Biology, 9(3), Daniel Hanus and Josep Call, Chimpanzees infer the location of a reward on the basis of the effect of its weight, pp. R370-372, © (2008), with permission from Elsevier.

As we have seen, it can be very informative to confront non-human subjects with situations that are perceptually very similar yet differ crucially in their conceptual content. The logic behind this approach is that if subjects go beyond perceptual information and encode the causally relevant structural properties of an object in a given context, they would be expected to behave differently in situations in which the same perceptual cue is related to the outcome in either a causal or a purely arbitrary way.

Hanus and Call (in press) conducted a second study to investigate chimpanzees' notion of weight. As in the "balance task" described earlier, arbitrarily and causally structured settings were contrasted but this time subjects not only observed events, but could also manipulate the task elements as it would usually appear under natural conditions. In the so-called "bottle task" chimpanzees were required to find and open the bottle containing fruit juice out of five possible opaque bottles of equal shape and size. Crucially, opening the bottles was costly in terms of the time and effort invested, thus emulating an active foraging situation. In the weight condition, all five bottles looked identical, but the juice bottle was much heavier than the other four empty bottles (causal cue). In the color condition, the weight of all five bottles was the same, since the other four bottles were filled with the same amount of water; the juice bottle, however, was color marked differently (arbitrary cue). Whereas in the weight condition, subjects rapidly learned to open the heavy bottle before the light bottles, the same individuals were not able to associate the conspicuous color of the bottle with its content, which resulted in a random opening order of the bottles. It is important to recall that in the given scenario, both types of information—the causal weight cue and the arbitrary color cue—possessed the same predictive value within the experiment, namely 100% in both cases. Nevertheless, for chimpanzees it seemed more intuitive to infer the content of an object based on structural causal information (in this case, heavy = food) than when given a purely perceptual cue based on an arbitrary regularity (in this case, white = food). In this experiment, the cues differed both in the nature of their relationship to the reward (causal versus arbitrary) and in their superficial characteristics. It remains possible that chimpanzees simply find weight more salient, because of either previous experience or a predisposition to attend to an object's weight rather than its color.

However, a recent study makes this explanation unlikely, because subjects were not more likely to learn weight than color cues in a discrimination task even though both cues held an arbitrary (but 100% predictive) relation to the reward (Schrauf & Call, 2009). In fact, some data suggested that the color cue was learned faster then weight, something that confirms previous studies showing that apes find discrimination based on weight particularly difficult to acquire (e.g., McCulloch, 1941). Thus, it is not the physical feature alone that induces the differences observed in these studies. Instead, we suggest that the context determines the relevance of the given information and specifies which cues are given precedence. Unlike Hanus and Call (in press), the color and weight cues in the Schrauf and Call (2009) setup are both arbitrarily linked to the outcome. This is likely to explain why the weight and color cues were learned equally quickly by the tested apes: neither of the cues was of a deterministic causal nature; hence, there was no reason to find one cue easier than the other.

As mentioned earlier, appreciating structural information is clearly beneficial because it enables an individual to generate reliable predictions without the need of extensive experience in a specific context. Furthermore, recognizing the causal relevance of a given property rather than having a pre-disposition to find certain properties more salient (such as preferentially attending to an object’s weight rather than its color) would be a more powerful strategy, given that in some contexts an object’s shape, for example, may be more important than its weight, but the reverse may be true in another context. Nevertheless, some properties may differ in the nature of their causal relation to reward. For instance, both color and weight are two potential indicators of the presence of an edible kernel inside a nut. However, whereas weight has a deterministic causal relation to the presence of a kernel—a light nut cannot contain an intact kernel—the right color and the presence of a kernel are not in the same sense causally related, because the nut might have been emptied by some insect. In other words, the necessity of a weight cue emerges from the fact that it might be unlikely in a natural setting for a nut with an “unripe color” to contain a ripe kernel, whereas it is simply impossible that a light nut contains a ripe kernel. Recent work has shown that subjects can select an appropriate object based on both weight and color cues. As well as the study by Schrauf and Call described earlier, wild capuchin monkeys have been shown to choose an appropriate stone hammer based not only on its visual appearance (size and color), but also its weight when these cues are removed or confounded (Visalberghi et al., 2009). Interestingly when the crucial information (e.g., weight) could not be judged by visual attributes (e.g., size), capuchins actively searched for additional information by touching and handling different stone alternatives. However, it remains to be seen whether non-human animals are sensitive to the difference between probabilistic and deterministic causal relations. A study investigating reactions to violations of different types of causal relationships would be an interesting topic for future study.

The findings on location problems fit well with those on extraction problems reviewed earlier and confirm and strengthen two key ideas. First, apes make a distinction between causal and arbitrary relations between stimuli. The causal-arbitrary distinction is a robust phenomenon found in a variety of setups including action- and perception-based measures. Second, there is no evidence of learning to solve the problem via conditional discrimination within the duration of the experiment—with enough trials they probably could but without such training they engage an inferential rather than associative mechanism, based on previously acquired knowledge of object properties.

(p.103) 3. Three kinds of knowledge about object properties

These recent experiments suggest that animals as well as young children are able to reinterpret first-order perceptual relations in terms of their higher order structural properties. Table 5.1 illustrates the important ways in which we think this abstract, *structural* knowledge differs from shallow, *perceptual* knowledge on one side, and from sophisticated *symbolic* knowledge on the other side.

As explained in Seed and Call (2009), the trap box can be used to illustrate the different sorts of information that can be extracted from the same external cue. In order to be successful, subjects need to move the reward away from the trap, but this solution could be based on perceptual, structural, or symbolic knowledge of the discriminatory stimuli:

- **Perceptual knowledge** consists of an arbitrary connection between the appearance of the box on any given trial and the correct response. Functionally relevant properties of the trap, such as its solidity, are not encoded. Generalization of the solution would therefore be limited to contexts that share some perceptual features with the initial task, and should be unaffected if the cue was made causally illogical (e.g., if the horizontal line was positioned above the shelf, made too small to trap the food, or made of a flimsy material). Furthermore, generalization would be limited to the modality in which the cue was originally encountered. For example, a subject reliant on perceptual knowledge based on the visual appearance of the task would not be expected to be able to solve the problem in the dark via tactile exploration.

- **Structural knowledge** would encompass some or all of the functional properties of the cue pertinent to its role in trapping the food, such as its solidity, or its continuity of surface. Generalization of the solution should therefore be possible across various changes to the perceptual elements of the original input, as long as the causal logic was unchanged. For instance, changing the color or the texture of a barrier should have little effect on performance. Furthermore, generalization to a perceptually similar task should be impaired if the task somehow violates causal logic, because the previously acquired structural knowledge should interfere. Another important difference is that structural knowledge should be multi-modal so that when information in one perceptual modality is not available, other senses can supply the required input to solve a task. For instance, if the subject has learned to move the reward away from a visible solid barrier, they should also be capable of using tactile or perhaps even auditory cues to locate that barrier and move the reward away from it.

- **Symbolic knowledge** cannot be directly re-interpreted from perceptual input in the way that structural knowledge can. Interestingly, in some of its characteristics it therefore resembles perceptual knowledge, because it consists of an arbitrary link between the cue and its referent. For example, the opaque version of the trap box task could be solved by moving the reward away from the Velcro cue that corresponds to the position of the trap, via perceptual knowledge or symbolic knowledge. However, in other ways these two types of knowledge differ dramatically. Whereas perceptual knowledge is based on first-order representations of the sort referred to by Mandler (2004) and others as "percepts," symbolic knowledge is (p.104) (p.105)

Table 5.1 Types of knowledge about causal properties

Knowledge type	Degree of abstraction	Nature of information	Generalization across stimuli that share …	Input modality	Information encoded as:	Example: Trap as a …
Perceptual	None	Arbitrary	… perceptual features	Unimodal	Percepts	… horizontal line
Structural	Abstract	Functional, Causal/Logical	… abstract, structural features	Multi-modal	Proto-concepts	… solid barrier
Symbolic	Abstract	Arbitrary	… abstract, conceptual features	Amodal	Concepts	… blue line cue standing for a solid barrier

> abstract and conceptual. A solution based on a symbolic cue could be transferred to a wide range of new contexts, including not only those that have no perceptual features in common with the trap box, but potentially also those that share no structural features. Whereas structural information is simply not encoded at the perceptual level, symbolic knowledge of a cue requires the subject to explicitly ignore its structural properties, to achieve what DeLoache refers to as dual representation. This provides a means for perceptual and symbolic knowledge to be teased apart, because whilst increasing the salience of a cue should facilitate learning at the perceptual level, conversely, it is likely to interfere at the symbolic level, because increased appreciation of the object itself may block an appreciation of its symbolic role (DeLoache, 2004).

Although the canonical example of symbolic representation is one that only bears an arbitrary relation with its referent, one type of symbolic representation deserves special mention. There are some cases in which the cue is not simply arbitrary because it bears an iconic relation with the referent. The Seed et al. connectedness problems to which we alluded to earlier fall within this category. Note that despite the scaffolding provided by the iconicity of the cue, subjects still did not use them to solve the task, something that they did when the cue possessed a causal relation with the reward.

This framework differs from the argument put forward by Penn, Holyoak, and Povinelli (Penn et al., 2008; Penn & Povinelli, 2007; Povinelli, 2000) in that we dispute their conclusion that the key difference between humans and other animals lies in an ability to extract meaningful, relational, abstract information from perceived stimuli. Although we agree that humans may be unique in their ability to use symbolic, abstract concepts to draw equivalence classes not grounded in or tied to causal logic, we think that the additional distinction we make between perceptual and structural information is worth highlighting. The ability to attribute inviolable causal properties such as weight and solidity to objects, or perhaps the degree to which individuals use this more abstract structural information to solve new problems, may be a difference between animals of different taxa not solely between humans and non-humans. For example, the failure of pigeons to solve the transfer tests in the connectedness task suggests that these birds may be particularly reliant on perceptual information. Indeed, the notion pigeons may use qualitatively different information to solve problems when compared to large-brained birds such as corvids has already been suggested by Wilson et al. (1985). They found in a number of paradigms that although corvids and pigeons learn equally quickly to make one response when stimuli match and another if they differ, corvids, but not pigeons, transfer this solution when a new set of stimuli are used, suggesting the use of relational rules by corvids and perceptual rules by pigeons.

Could differences in the salience of different perceptual cues within different contexts explain the differences we find between the causal and arbitrary conditions in the paradigms we use? Such an account is hard to falsify, but we argue that apes and corvids use the causal context to distinguish which facet of an object is relevant in a given context (e.g., its weight or its shape). This is because some subjects can solve problems such as the two-trap transfers in which the same cue must be responded to differently in different contexts (a solid object can be a barrier or a supporting surface). Another striking example is the fact that chimpanzees treated the downward movement of the balance beam differently, depending on whether an external cause acted on it. In this task in particular, but probably also in others that we have described, knowledge of object properties must interact with knowledge of causal interrelations which goes beyond pure spatiotemporal associations between events. In this example, it is not enough to know that an object has intrinsic weight, one must also recognize when this property will have causal relevance. Knowledge of object properties would therefore be much more beneficial if they are embedded in a causal map that dictates which facet of an object is relevant in a given causal structure. An ability to infer causal directionality—for example, from conditional probability—would be one route to the acquisition of causal knowledge (if all round, heavy objects, and all brown, heavy objects, but no light, round, brown objects make good hammers, one can infer that heaviness, but not shape or color is a causal property in this context). Causal learning certainly plays an important part in children’s attribution of causal properties (Schultz et al., 2007), but the extent to which it plays a role in animal causal judgments is yet to be explored.

One potentially important finding arising from the research presented in this chapter is the concordance between extraction problems and location problems. The evidence is not massive to date but there are some good indications. Subjects were able to solve problems that required an appreciation of the effects that a trap has on a moving reward both when using tools to get the reward and when they needed to infer the location of the reward after an invisible displacement. With regard to weight, subjects were able to locate the reward both by feeling the size of identically looking containers and by using the effect that the reward had on a balance beam. Although it is conceivable that subjects solve this tasks independently from each other, as if each were encapsulated in its own domain, it is also possible that subjects do indeed possess a more general conception of knowledge about traps/obstacles and weights. One hypothesis is that subjects possess knowledge about traps and weight that allowed them to respond to various challenges ranging from extracting visible food while avoiding obstacles to inferring the location of hidden reward. Structural knowledge, unlike perceptual knowledge that would be tied to particular stimuli, would be a good candidate for encoding the information used in solving these tasks.

What does it mean if a subject does not recognize the causal relevance of a property in one setup, when the results of a previous task indicated they did use knowledge of this same property? One possibility is that there are limits to the connections that individuals may make between problems in different domains. It is conceivable that although subjects use the same substrate of knowledge to solve different tasks with different demands, they do not explicitly recognize the relation that exists between the various elements in those tasks. In other words, subjects may be aware of the causal relations between the elements within a particular problem, but they do not establish analogical relations between functionally common elements across tasks. It is also important to recognize that even slight changes in the motor or perceptual demands within each task can have serious detrimental consequences for performance as previously indicated. The three types of knowledge described in our table should not be seen as mutually exclusive. Subjects may rely on a perceptual strategy to learn a response to the discriminative cue when certain task demands (p.107) are increased, for instance by changing the perceptual or motor task demands. This could mean that knowledge that subjects may possess about causal relations is not as robust as one may find in human adults, whose actions are nevertheless also influenced by perceptual as well as conceptual features of physical problems (Silva et al., 2005, 2008).

4. Future directions

Indeed, there is much still to learn about how non-human animals understand causality, perhaps because the results of early experiments suggested that they do not cognize causal structures or causal properties. Recent experiments cast doubt on this assumption, and we suggest that previous negative results may be in part ascribable to limitations in other psychological faculties, such as inhibition and working memory, especially in tests requiring the use of tools. Our hypothesis is that rather than there being a great divide between humans and animals, with only humans reinterpreting the world around them in terms of higher order properties and causal structures, there may be some species that also form abstract, multi-modal representations encompassing some of the structural properties of objects. A number of questions arise from this theoretical position: which species, which specific properties, how do the abilities develop, how are they algorithmically and physically realized, and how do they differ from those of humans? With regard to the latter question we think that symbolic knowledge, not causal knowledge, may be the point of departure. Attributing a causal relationship between an arbitrary cue and an outcome, especially in contexts where the structural properties of the cue mean that attributing causal power to it would violate existing causal knowledge, may require additional cognitive machinery only available to humans.

References

Bibliography references:

Bitterman, M. E. (1975). The comparative analysis of learning. *Science*, **188**, 699–709.

Blaisdell, A. P., Sawa, K., Leising, K. J., & Waldmann, M. R. (2006). Causal reasoning in rats. *Science*, **311**, 1020–2.

Boysen, S. T., Berntson, G. G., Shreyer, T. A., & Quigley, K. S. (1993). Processing of ordinality and transitivity by chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology*, **107**, 208–15.

Call, J. (2001). Object permanence in orangutans (*Pongo pygmaeus*), chimpanzees (*Pan troglodytes*), and children (*Homo sapiens*). *Journal of Comparative Psychology*, **115**, 159–71.

Call, J. (2010). Trapping the minds of apes: causal knowledge and inferential reasoning about object-object interactions. In E. Lonsdorf, S. R. Ross, and T. Matsuzawa (eds), *The Mind of the Chimpanzee: Ecological and Experimental Perspectives* (pp. 75–86). Chicago/London: University of Chicago Press.

Call, J., & Tomasello, M. (2005) Reasoning and thinking in nonhuman primates. In K. J. Holyoak & R. G. Morrison (eds), *The Cambridge Handbook of Thinking and Reasoning* (pp. 607–32). New York: Cambridge University Press.

Cunningham, C. L., Anderson, J. R., & Mootnick, A. R. (2006). Object manipulation to obtain a food reward in hoolock gibbons, *Bunopithecus hoolock*. *Animal Behaviour*, **71**, 621–9.

de Blois, S. T., Novak, M. A., & Bond, M. (1998). Object permanence in orangutans (*Pongo pygmaeus*) and squirrel monkeys (*Saimiri sciureus*). *Journal of Comparative Psychology*, **112**, 137–52.

DeLoache, J. S. (2004). Becoming symbol-minded. *Trends in Cognitive Sciences*, **8**, 66–70.

Dickinson, A., & Shanks, D. (1995). Instrumental action and causal representation. In D. Sperber, D. Premack, & A.

Premack (eds), *Causal Cognition: a multidisciplinary debate* (pp. 5–25). Oxford: Oxford University Press.

Fujita, K., Kuroshima, H., & Asai, S. (2003). How do tufted capuchin monkeys (*Cebus apella*) understand causality involved in tool use? *Journal of Experimental Psychology Animal Behavior Processes*, **29**, 233–42.

Gillan, D. J. (1981). Reasoning in the chimpanzee: II. Transitive inference. *Journal of Experimental Psychology Animal Behavior Processes*, **7**(2), 150–65.

Girndt, A., Meier, T., & Call, J. (2008). Task constraints mask great apes' ability to solve the trap-table task. *Journal of Experimental Psychology Animal Behavior Processes*, **34**, 54–62.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, **111**, 3–32.

Hanus, D., & Call, J. (2008). Chimpanzees infer the location of a reward on the basis of the effect of its weight. *Current Biology*, **18**, R370–2.

Hauser, M. D., Kralik, J., & Botto-Mahan, C. (1999). Problem solving and functional design features: Experiments on cotton-top tamarins, *Saguinus oedipus oedipus*. *Animal Behaviour*, **57**, 565–82.

Hauser, M. D., Pearson, H. & Seelig, D. (2002). Ontogeny of tool use in cottontop tamarins, *Saguinus oedipus*: Innate recognition of functionally relevant features. *Animal Behaviour*, **64**, 299–311.

Hauser, M. D., & Spaulding, B. (2006). Wild rhesus monkeys generate causal inferences about possible and impossible physical transformations in the absence of experience. *Proceedings of the National Academy of Sciences of the USA* **103**, 7181–5.

Herrmann, E., Wobber, V., & Call, J. (2008). Great apes' (*Pan troglodytes*, *Pan paniscus*, *Gorilla gorilla*, *Pongo pygmaeus*) understanding of tool functional properties after limited experience. *Journal of Comparative Psychology*, **122**, 220–30.

Limongelli, L., Boysen, S. T., & Visalberghi, E. (1995). Comprehension of cause-effect relations in a tool-using task by chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology*, **109**, 18–26.

Mandler, J.,M. (2004) *The Foundations of Mind: Origins of Conceptual Thought.* New York: Oxford University Press.

Martin-Ordas, G., & Call, J. (2009). Assessing generalization within and between trap tasks in the great apes. *International Journal of Comparative Psychology*, **22**, 43–60.

Martin-Ordas, G., Call, J., & Colmenares, F. (2008). Tubes, tables and traps: Great apes solve two functionally equivalent trap tasks but show no evidence of transfer across tasks. *Animal Cognition*, **11**, 423–30.

McCulloch, (1941). Discrimination of lifted weights by chimpanzees. *Journal of Comparative Psychology*, **32**, 507–19.

Mendes, N., Rakoczy, H., & Call, J. (2008). Ape metaphysics: Object individuation without language. *Cognition*, **106**, 730–49.

Mulcahy, N. J., & Call, J. (2006). How great apes perform on a modified trap-tube task. *Animal Cognition*, **9**, 193–9.

Natale, F. (1989). Causality II: The stick problem. In F. Antinucci (ed.), *Cognitive Structure and Development in Nonhuman Primates* (pp. 121–33). Hillsdale, NJ: Lawrence Erlbaum Associates.

Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin’s mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, **31**, 109–30.

Penn, D. C., & Povinelli, D. J. (2007). Causal cognition in human and nonhuman animals: a comparative, critical review. *Annual Review of Psychology*, **58**, 97–118.

Phillips, W., & Santos, L. R. (2007). Evidence for kind representations in the absence of language: Experiments with rhesus monkeys (*Macaca mulatta*). *Cognition*, **102**, 455–63.

Piaget, J. (1952). *The Origins of Intelligence in Children*. Oxford: International Universities Press.

Povinelli, D. J. (2000). *Folk Physics for Apes: The Chimpanzee's Theory of How the World Works*. Oxford: Oxford University Press.

Premack, D. (1983). Animal cognition. *Annual Review of Psychology*, **34**, 351–362.

Santos, L. R., Sulkowski, G. M., Spaepen, G. M., & Hauser, M. D. (2002). Object individuation using property/kind information in rhesus macaques (*Macaca mulatta*). *Cognition*, **83**, 241–64.

Schmidt, G. F., & Cook, R. G. (2006). Mind the gap: Means-end discrimination by pigeons. *Animal Behaviour*, **71**, 599–608.

Schrauf, C., & Call, J. (2009). Great apes' performance in discriminating weight and achromatic color. *Animal Cognition*, **12**, 567–74.

Schulz, L. E., Kushnir, T., Gopnik, A. (2007) Learning from doing: intervention and causal inference. In A. Gopnik and L. E. Schultz (eds), *Causal Learning: Psychology, Philosophy and Computation* (pp. 67–85). Oxford, Oxford University Press.

Seed, A. M., Albiach-Serrano, A., & Call, J. (in preparation). Don't mind the gap: Bonobos, chimpanzees and children use functional but not arbitrary information to solve the broken string problem.

Seed, A. M., & Call, J. (2009). Causal knowledge for events and objects in animals. In S. Watanabe, A. P. Blaisdell, L. Huber & A. Young (eds), *Rational Animals, Irrational Humans* (pp. 173–87). Tokyo: Keio University Press.

Seed, A. M., & Call, J. (in preparation). Children use functional, not arbitrary information to solve the trap problem.

Seed, A. M., Tebbich, S., Emery, N. J., & Clayton, N. S. (2006). Investigating physical cognition in rooks, *Corvus frugilegus*. *Current Biology*, **16**, 697–701.

Silva, F. J., Silva, K. M., Cover, K. R., Leslie, A. M., & Rubalcaba, M. A. (2008). Humans' folk physics is sensitive to physical connection and contact between a tool and reward. *Behavioural Processes*, **77**, 327–33.

Silva, F. J., Page, D. M., & Silva, K. M. (2005). Methodological-conceptual problems on the study of chimpanzees' folk physics: How studies with adult humans can help. *Learning and Behaviour*, **33**, 47–58.

Spinozzi, G., & Potí, P. (1989). Causality I: The support problem. In F. Antinucci (ed.), *Cognitive Structure and Development in Nonhuman Primates* (pp. 114–19). Hillsdale, NJ: Lawrence Erlbaum Associates.

Taylor, A. H., Hunt, G. R., Medina, F. S., & Gray, R. D. (2009). Do New Caledonian crows solve physical problems through causal reasoning? *Proceedings of the Royal Society B: Biological Sciences*, **267**, 247–54.

Tebbich, S., & Bshary, R. (2004). Cognitive abilities related to tool use in the woodpecker finch, *Cactospiza pallida*. *Animal Behaviour*, **67**, 689–97.

Tebbich, S., Seed, A. M., Emery, N. J., & Clayton, N. S. (2007). Non-tool-using rooks, *Corvus frugilegus*, solve the trap-tube problem. *Animal Cognition*: **10**, 225–31.

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative process in animals. *Psychological Review, Monograph Supplements*, **2**, 109.

Tomasello, M., & Call, J. (1997). *Primate Cognition*. New York: Oxford University Press.

Visalberghi, E., Addessi, E., Truppa, V., Spagnoletti, N., Ottoni, E., Izar, P., & Fragaszy, D. (2009). Selection of effective stone tools by wild bearded capuchin monkeys. *Current Biology*, **19**, 213–17.

Visalberghi, E., & Limongelli, L. (1994). Lack of comprehension of cause-effect relations in tool-using capuchin monkeys (*Cebus apella*). *Journal of Comparative Psychology*, **108**, 15–22.

Waldmann, M. R., Hagmayer, Y., & Blaisdell, A. P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science*, **15**, 307–11.

Wilson, B., Mackintosh, N. J., & Boakes, R. A. (1985). Transfer of relational rules in matching and oddity learning by pigeons and corvids. *Quarterly Journal of Experimental Psychology*, **B 37**, 313–32.