Chapter 9

The rationality of animal memory: Complex caching strategies of western scrub jays

Nicky Clayton, Nathan Emery, and Anthony Dickinson

Abstract

Scrub jays cache perishable and non-perishable foods, and their caches may be pilfered by conspecifics. Caching and recovery by scrub jays is psychologically rational in the sense that these behaviours responded appropriately to conditions that should have changed the birds' beliefs and desires. For example scrub jays were allowed to cache worms and peanuts in a visuospatially distinct tray. At recovery, birds search initially for worms after a short retention interval because they believe that the worms are still edible, but switch to searching for peanuts at a long retention interval because they believe that worms are now degraded. If jays acquire new information after caching, such that worms are no longer edible when recovered at the short interval, this should affect their belief about the state of their caches. Jays update their cache memory, and on subsequent trials of the short interval, search selectively in peanut sites. In a second example, scrub jays cached either in private (when another bird's view was obscured) or while a conspecific was watching, and then recovered their caches in private. Scrub jays with prior experience of stealing another bird's caches subsequently recached food in new sites during recovery trials, but only when they had been observed caching. Naïve birds did not. We suggest that experienced pilferers had formed a belief that observers will pilfer caches they have seen, and recache food in new sites to fulfil their desire to protect their caches. Since recaching is not dependent on the presence of the potential thief, the jays must recall the previous social context during caching, and flexibly use this information to implement an appropriate cache protection strategy, namely recache the food in locations unbeknownst to the pilferer.

9.1 Introduction: intentional and mechanistic psychology

In common with other scatter-hoarding animals, western scrub jays (*Aphelocoma californica*) hide surplus food in discrete locations within their territories, which they recover in times of need (Van der Wall 1990). Psychology offers two classes of explanation for such behaviour: the mechanistic and the intentional. Mechanistic accounts appeal to psychological processes that gain their explanatory power by analogy to physical processes. A classic example is associative learning theory. According to this theory, learning about the relationship between events, whether they are stimuli or responses, consists of the formation of excitatory (or inhibitory) connections between nodes activated by these events. The mechanism by which such associative structures control behaviour is the transmission of excitation (or inhibition) from one node to another until the activation of the terminal node of the associative chain is sufficient to generate the observed behaviour.

Associative theories have had an enduring influence on the study of animal learning and cognition ever since Thorndike (Thorndike 1911) formulated his Law of Effect on the basis of the first controlled, comparative studies of learning in animals. The development of associative theory continued throughout the last century to reach the complexity and sophistication of multilayered networks with distributed representations (Rumelhart and McClelland 1986). We do not need to appeal to such complex structures, however, to explain cache recovery by jays in terms of associative mechanisms. All that we need to assume is that associative nodes activated by the visual cues around the cache become connected with those excited by the food stored at that site, perhaps a peanut, at the time of caching. Consequently, re-exposure to the cache site stimuli activates the food nodes that, in turn, are associated with nodes controlling an approach response either innately or through prior learning.

We suspect, however, that we should have little success in persuading our proverbial grandmother of the merits of this account. By analogy with remembering where she hid her cache of chocolate when overtaken by a craving, our granny would probably explain the jay's recovery behaviour in terms of beliefs and desires. Her explanation might run something like this: the jay, being hungry, has a desire for food and, as a result of caching the peanut, has a belief that searching in the cache site will yield food. If pressed about why having this belief and desire causes the bird to search the cache site, granny would probably tell us to stop being obtuse—given this belief and desire, the only sensible or rational thing to do is to search the cache site. Moreover, if she would tolerate further Socratic enquiry, we should find that what she means is that cache searching is rational because, of necessity, it must fulfil the jay's desire for food if the belief that searching the cache site will yield food is true.

It is the rationality of this practical inference process that distinguishes the intentional explanation from the mechanistic one. Our account of rational behaviour accords with that outlined by Dretske (Chapter 3, this volume). The processes by which

associative structures control behaviour are constrained only by their mechanistic-like properties. The level of activation of the food nodes, and hence the probability and vigour of searching, is simply a function of the strength of the associative connections, the input activation, the nodes' thresholds for activation, whether or not there is concurrent inhibitory input, etc. But in and of themselves, none of these processes yield behaviour that necessarily conforms to any cannons or principles of rationality. Of course, evolution will have ensured that behaviour governed by such mechanistic processes is adaptive in that it contributes to the reproductive fitness of the jay by maintaining its nutritional state, and in this sense associatively-controlled cache searching can be regarded as biologically rational (see Kacelnik, Chapter 2, this volume). But such biological rationality must be distinguished from the psychological rationality of the practical inference process that generates cache searching from a belief about the consequences of this behaviour and the desire for these consequences.

According to this analysis, the issue of whether an animal is psychologically rational turns on the nature of the processes causing its behaviour; specifically on whether this behaviour is caused by psychological mechanisms or by intentional processes. The jay's behaviour is psychologically rational to the extent that it is caused by the interaction of a belief and desire in such a way that performance of the behaviour in question fulfils the desire if the belief is true (and fails to do so if the belief is false). Such an account is intentional because it requires that the antecedent mental states, the belief and the desire, have intentional properties, such as truth and fulfilment, because their content represents current or desired states of affairs.

There are a number of points to note about this analysis. First, it is not sufficient for an intentional explanation that the behaviour is simply caused by a belief and desire; rather that it has to be caused by the right process, namely the rational process of practical inference. It may well be that the sight of a cache site in which a food-desiring jay believes that searching will yield food causes an increase in heart rate as a component of general autonomic arousal. But such activation is not caused intentionally. Unlike cache searching, whether or not the jay's autonomic nervous system is activated has no necessary relation to whether its desire will be fulfilled and is, therefore, to be explained in terms of psychological mechanisms rather than by intentional processes. Perhaps in the past being in these belief and desire states has been associated with autonomic arousal so that the reinstatement of these states once again triggers an increase in heart rate.

¹ Editors' note: For related distinctions, see Kacelnik on PP rationality and Dretske on minimal rationality; cf. Allen, Papineau and Heyes for methodological reflections and scepticism about related distinctions (all in this volume). Dretske distinguishes behaviour that is (merely) caused by a state with intentional properties, and the content of a state with intentional properties being explanatorily relevant to some behaviour; on his view, the latter is a necessary condition for rational behaviour.

Second, according to this analysis, rationality is a property not of an animal, but of the processes causing its behaviour. So, to characterize a jay as rational does not imply that all of its behaviour is intentional, but rather that it is capable of at least some intentional action. We do, of course, take humans as the canonical case of a rational animal because the concordance between our actions and the expression of our beliefs and desire through language provide evidence of intentional causation. This is not to say, however, that all of our behaviour is rational. In fact, we may well be surprised to discover how little of our daily life is in fact under intentional control! The complex sequences of behaviour by which we drive, walk, or cycle to work may well appear to be purposive and goal-directed but, on further investigation, to consist of a chain of mechanistically elicited habits, albeit complex and highly structured ones, triggered by the stimuli along our route. So even we are creatures of a dual psychology; the mechanistic and the intentional. Therefore the issue of whether an animal is psychologically rational is really the issue of whether any of its behaviour warrants an intentional account.

Finally, it is important to note that behaviour does not carry its rationality on its sleeve. There may be nothing obvious about the manifest behaviour during our journey to work that marks whether it is under intentional or habitual control. Indeed, it is most likely that it started out under intentional control and only became habitual with repetition (Dickinson 1985, 1989). And, for the same reason, the processes of behavioural control cannot be determined by simple observation of the adaptive nature of the behaviour (Heyes and Dickinson 1990). As we shall see in our discussion of the psychological processes controlling food caching and recovery by western scrub jays, determining whether these processes are *psychologically* mechanistic or intentional is a complex enterprise which requires converging lines of evidence.

9.2 The content of desires

For a variety of reasons, there is a reluctance to accept that non-linguistic creatures can be endowed with an intentional and therefore a rational psychology. One issue relates to the vehicles that carry the content of beliefs and desires and, specifically, to whether only explicit and therefore potentially public languages, be they natural or artificial, are capable of providing the requisite vehicles for the content of beliefs and desires. This is not the place to debate this issue other than to note that certain philosophers of mind, most notably Fodor (1977), have vigorously argued that explicit languages are grounded on a language of thought that is shared, at least in some rudimentary form, by some non-linguistic creatures. By contrast, others have disputed the claim that a language, natural or purely mental, provides the only psychological vehicles for representational or intentional content (see Bermúdez, Chapter 5, this volume).

If one accepts, at least in principle, that animals other than humans can possess mental states with intentional or representational content, there still remains the problem of

determining their content in the absence of an explicit, communicable language. As we have already noted, unlike speech acts, behaviour does not manifest its intentionality explicitly. However, it is possible to make some empirical progress on this issue in specific cases.

Consider once again our jay recovering a peanut from its cache site. Is this behaviour motivated by a desire for food or by a desire for a peanut? In other words, is the bird motivated by a general desire or by a specific one? We have attempted to answer this question by using the technique of devaluing the food immediately prior to recovery. In this study (Clayton and Dickinson 1999), jays cached peanuts in one cache site and dog food kibbles in a second cache site before being allowed them to search for these foods after a retention interval. The important feature of this study was that immediately prior to recovery we attempted to remove the desire for one of the foods, while maintaining the desire for the other. It is well known that we and other animals show food-specific satiety—a surfeit of one type of food, however delicious initially, rapidly looses its pleasure while maintaining, or even enhancing, the attraction of other foods. Consequently, the birds were prefed one of the two foods, either peanuts or kibbles, immediately prior to giving them a choice between searching in the two cache sites. To the extent that cache searching was motivated by a general desire for food, this prefeeding should have produced just a general decrease in their desire for food and hence an equivalent reduction in searching in both sites. By contrast, motivation by a specific desire for peanuts should have reduced searching selectively in the site in which the birds had cached the kibbles, and the opposite searching preference should be shown by jays that had been prefed peanuts.

The results favoured the selective rather than the general content of desires. The birds searched preferentially in the cache sites in which they had cached the non-prefed food; in other words, the birds that had been prefed peanuts searched in the kibbles site and those that had been prefed kibbles searched in the peanut site. This finding corresponds with what we know about motivation in other animals for which food-related desires are both specific and learned (Dickinson and Balleine 1994, 2002). An important feature of the experimental design is that we pilfered the food from the cache sites before allowing the birds to search for their caches on test trials. If we had not done so, as soon as the birds had recovered a cache of the prefed type, they would have discovered that this food was no longer attractive and thereby extinguished searching in this site, and, as a consequence, their pattern of searching would not have needed to reflect the interaction of their relative desires for the two foods with beliefs about where they were cached. By pilfering the caches we ensured that on test trials the relative preference must have been mediated by memory for the location of these caches.

What this procedure did not ensure, however, is that this memory took the form of a belief rather than of an association between the cues of the cache site and the type of food stored there (or, more strictly speaking, between the nodes activated by the cues and food). An associative account of the preference only requires that the repetitive

activation of the food nodes during pre-feeding induced a temporary refractory state in these nodes. Consequently, the activation of the prefed food nodes at the time of recovery by the cache-site cues would have been reduced by their refractory states, thereby producing a preference for searching in the cache site of the non-prefed food. In an attempt to differentiate between the associative and the intentional accounts, we repeated the specific-satiety procedure using the following, more complex, design.

To understand this design, it is necessary to describe our procedures in more detail. The cache sites used in our studies are ice-cube trays, each consisting of two parallel rows of ice-cube moulds filled with a substrate such as sand. The trays are made spatially and visually distinct by surrounding each of them with a structure built of toy building blocks of various colours and shapes so that they are topographically unique. In this way, a given bird caches in different trays on different days, without having to reuse the trays. The birds must attend to and learn about the trays' cues because they cache different foods in the two sides of a particular tray and have to use these cues to remember where in a tray they have cached a particular type of food.

So it was in the experiment outlined in Fig. 9.1. Each bird cached three peanuts in one side of two trays, the Same and Different Trays, on separate caching episodes, and three kibbles in the other side of each tray on two further caching episodes. For example the jay, whose design is illustrated in Fig. 9.1, cached the peanuts in the left side of the trays during caching episodes 1 and 2, and the kibbles in the right sides of the trays during caching episodes 3 and 4. On the assumption that each pairing of the




Fig. 9.1 The experimental design for differentiating between the associative and the intentional account of cache recovery. Numbers in italics represent the increase in associative strength of the connection between the tray cue and peanut (p) and kibble (k) nodes. When restricted to caching or recovering a particular food type in just one side of the caching tray, the other side of the tray is covered with a Plexiglas cover. The Plexiglas strip is removed during test so that birds can search both sides of the caching tray. P: peanuts; K: kibbles. The brackets indicate the caches that would have remained in the tray at recovery if they had not been previously removed on test.

tray cues with a food item enhanced the associative strength of the connection between the tray cue nodes and the food nodes, the strength of the connection between the tray cue nodes and the peanut nodes should have been 3p where p is the increment in associative strength produced by a single peanut. Similarly, tray cue-kibble associative strength should have been 3k where k is the increment in strength produced by pairing a single kibble with the trays' cues. Therefore, after these caching episodes both the same and the different trays cues should have a total associative strength with food nodes of 3p + 3k (Fig. 9.1). The trays were then removed from the birds' homecages for a retention interval of 3 hours.

After the retention interval, the trays were returned to the birds who were then allowed to recover the three peanuts from the Same Tray and the three kibbles from the Different Tray. These tray cue—food pairings should have further enhanced the total associative strengths to 6p + 3k for the Same Tray but to 3p + 6k for the Different Tray. Finally, the birds were prefed one of the foods to satiety before being given the opportunity to search again in the two trays, but this time with both sides of each tray available. At issue is whether they directed most searches to the Same Tray that should have still contained the food of the *same* type as that which had been prefed, or to the Different Tray that should have still contained the food that was *different* to the prefed type. Of course, we had pilfered all the remaining food items prior to this final recovery test.

The prediction of the associative theory is clear: the birds should have searched preferentially in the Same Tray. Recall that prior to prefeeding the Same Tray had a total associative strength of 6p + 3k whereas the Different Tray had a strength of 3p + 6k. On the assumption that prefeeding renders the nodes for the same food refractory, the total activation of the food nodes by the tray cues should have been 6p for the Same Tray but only 3p for the Different Tray after prefeeding on kibbles. If searching in a tray is determined by the strength of activation of food nodes by the tray cues, then the jays should have searched in the same tray rather than the different one.

But clearly, such a search pattern would not be the rational choice. After the regime of caching, recovery, and prefeeding, the Same Tray should have contained only the less desirable food items, the kibbles in the example illustrated in Fig. 9.1, whereas the desirable food, namely peanuts, should have been in the Different Tray. Our birds behaved rationally in this test by showing a marked preference for searching for what should have been intact caches in the Different Tray. This finding accords with an intentional account that assumes that searching is controlled by specific desires interacting with beliefs about the content of the cache sites, beliefs that were acquired during caching but were also updated by subsequent experience.

There are two general points to be drawn from these studies. The first is that, given one adopts an intentional account of animal behaviour, the problem of determining the content of intentional representations is not an entirely intractable one. Second, these studies illustrate an empirical strategy for evaluating intentional, and therefore

psychologically rational, accounts of behaviour. There is never going to be a behaviourally decisive test of the intentionality of animal behaviour. All one can do is to derive predictions from an intentional account of a specific behaviour, in this case cache recovery, and test these predictions within a procedure that discriminates this explanation from a specific, mechanistic alternative such as the associative account (see and cf. Papineau and Heyes, Chapter 8, this volume; see also Kacelnik, Chapter 2, this volume, for scepticism about the empirical tractability of 'PP rationality' in studies of animals). Having illustrated how this strategy can be implemented in the case of desires, we now turn to the investigation of cache beliefs.

9.3 The structure and content of cache beliefs

Human cognitive psychology classifies beliefs as declarative memories to distinguish them from procedural memories, which encompass various forms of acquired motor and cognitive skills, responses, and habits. Declarative memory can be further subdivided into two forms: semantic or general knowledge and episodic recall (Tulving 1972, 1983). Semantic memories are general beliefs or knowledge, whereas episodic memories are recollections of specific, particular life events. Thus, for a jay the knowledge about the location and properties of the reliable sources of food in its territory would be an example of general knowledge, whereas the recollection of a particular caching episode at the time of recovery would be an example of episodic memory. Contemporary accounts of human memory (i.e. Tulving and Markowitch 1998) view episodic memory as being embedded within a more general declarative framework in such a way that specific episodic information can interact with general declarative knowledge. This interaction ensures that action based upon information represented in a recollection of specific episode is informed by the agent's corpus of general knowledge.

For a number of years, we have been investigating whether cache recovery by western scrub jays is mediated by at least a declarative-like memory (Clayton and Griffiths 2002; Clayton *et al.* 2000; Clayton *et al.* 2001; Griffiths *et al.* 1999). Our studies capitalized on the fact that these jays are omnivorous, eating and caching a variety of foods such as insects, larvae, and nuts. One problem the jays face, living as they do in the Californian Central Valley, is that some of their most preferred foods, such as invertebrates, decay if left too long in the cache before recovery. In contrast, other foods, such as nuts, are relatively durable. Consequently, as a result of experience with caching and recovering various foods after different cache-recovery (retention) intervals, the jays may acquire general knowledge about the rates at which different foods perish which they can then deploy in conjunction with episodic-like memory for specific caching events to determine their choices at recovery.

As a concrete illustration of this interaction between general knowledge and episodic-like recall, consider the case in which a jay caches perishable crickets in one site and non-perishable peanuts in another. As crickets are one of the bird's preferred

foods, it should choose to search for cricket caches in preference to peanut caches. The problem is, however, that crickets perish if left in the cache for too long, so the choice at recovery should depend upon the length of the retention interval. When recovering caches after a short retention interval, the jay should search preferentially for crickets, but this preference should reverse after longer retention intervals so that the bird searches for peanuts if it believes that the crickets will have perished.

The fact that our colony of scrub jays are hand-raised, and therefore have no prior experience with decaying foods, allowed us (Clayton *et al.* 2001) to investigate whether the birds are capable of learning this reversal in recovery preference. As in the previous study, on each trial the birds cached one food in one side of a trial-unique caching tray and the second food in the other side of the same tray, but in this case the foods were crickets and peanuts. On some trials they were allowed to recover both peanuts and fresh crickets after one day, whereas on other trials the opportunity for recovery was delayed for 4 days by which time the crickets had decayed.²

The jays rapidly learned to search for crickets when fresh and to search for peanuts when the time interval between caching and recovery was such that the crickets should have degraded. On the first two trials the majority of birds directed their first search to the cricket side of the tray after the 4-day retention interval, but by the third trial all birds switched their preference and searched in the peanut side first. However, this preference switch was under temporal control because the majority of birds continued to direct their first search to the cricket side on all training with the 1-day retention interval. The reason why we recorded the first direction of the first search on these training trials was because once a bird had found a food cache it did not have to rely on memory for its caches to determine its preference—all it need do was avoid the side on which it had just found a less preferred food item, be it a peanut or decayed cricket.

It remains possible, however, that the birds could have detected the type of food buried in each side of the caching trays before making even their first search, and so we conducted a series of probe trials in which we pilfered all the food caches prior to the recovery test. In addition, we also tested searching at recovery after untrained retention intervals of 2, 3, and 5 days to see how the birds interpolated and extrapolated from the trained retention intervals of 1 and 4 days (Clayton *et al.* 2003). The profile of searching shows that the jays had inferred from their training experience that crickets remained palatable for up to 3 days but then perished. The majority of birds searched the cricket side first after retention intervals of up to 3 days and before switching their preference to the peanut side after longer intervals.

² Clayton *et al.* (2001) included trials with a third food type, wax moth larvae, and included a shorter 4-hour retention interval, but the procedural description has been simplified for exposition in a way that does not vitiate the interpretation.




Fig. 9.2 The architecture and contents of a jay's declarative memory (see text).

The ability to learn about the degradation of crickets and to apply this knowledge to new cache sites is readily analysable in terms of declarative memory. A possible architecture for the jays' declarative memory in this task is illustrated in Fig. 9.2. According to this model, the jays acquired two forms of general knowledge during their training experience. First, they learned about the relative desirability of the various foods they experienced: fresh crickets, degraded crickets and peanuts. Second, they acquire beliefs about the temporal decay of the crickets: specifically that crickets are fresh at retention intervals (RIs) of less than 4 days, but degraded after longer intervals.

The second component of the declarative system is an episodic-like memory which encodes in a bound and integrated representation the content of a specific life event, which in this case is the caching of peanuts in the left side and crickets in the right side of a particular tray 3 days ago. We refer to such a representation as a 'what—where—when' memory because it is important that all three features of the experience are represented if this memory is to resemble episodic memory. Recent discussions of human episodic memory have emphasized the phenomenological characteristics of recollection, such as 'autonoetic awareness' (Tulving and Markowitsch 1998). However, such awareness cannot be assessed in non-linguistic animals because there are no agreed behavioural markers of consciousness in the absence of language. It is for this reason that we have referred to the memories mediating cache recovery by western scrub jays as episodic-like (Clayton *et al.* 2000, 2001; Griffiths *et al.* 1999) by reference to Tulving's (1972) original characterization of episodic memory as a form of memory that 'receives and stores information about temporally dated episodes or events, and temporal–spatial

relations among those events' (p. 385). By this 'what-where-when' criterion, the bird's memory of the caching episodes is episodic-like because it involves recall of the content (what) and location (where) of their cache, and a temporal component (when).

This model allows us to offer a sketch of the interaction between general knowledge and episodic-like memories through the processes of practical inference. Our assumption is that the birds have an enduring intention to search for the most desired food items when in the motivational state that promotes cache recovery. Encountering a cache site, in this case a particular tray, retrieves the what-where-when episodiclike memory of caching in that tray. Integrating this memory with general knowledge about its relative food desires and the degradation profiles of these foods allows the bird to derive the more specific intention to search for crickets. This intention, when taken in conjunction with the episodic-like memory that crickets were cached in right-hand side of the tray, leads to the derivation of an intention that can be directly expressed in behaviour, specifically the intention to search in the right-hand side.

9.3.1 The flexibility of cache memories

We do not intend that the details of this model should be taken at face value. Rather its function is to illustrate a cardinal feature of declarative memory systems that is central to the issue of animal rationality, namely that declarative memories or beliefs have a representational form that allows their content to be deployed flexibly. It was this issue of the flexibility of knowledge representations that motivated Winograd's (1975) classic analysis of the relative merits of declarative and procedural representations for artificial intelligence. In that analysis, he pointed out that the flexibility of declarative representations is bought at the computational cost of having inference processes that can operate on their content. Indeed Tulving (2001) explicitly raised the issue of mnemonic flexibility with reference to the role of declarative memory in cache recovery when he asked '... could Clayton and her colleagues (or someone else) get their scrub jays, who remember what kind of food is where, to do something other with that information than act on it 'inflexibly'...?' (p. 1513).

Although the concept of 'flexibility' is not well defined, our model of declarative memory allows for the flexible deployment of episodic-like information in relation to a bird's general, semantic-like, knowledge; a flexibility that arises from their interaction through practical inference processes. Recall that our jays experienced the variable palatability of crickets only after 1 and 4 days during training, so their belief that crickets degrade between 3 and 4 days after caching is a generalization from this training experience. But what would happen if we falsified this belief during the retention interval for a caching episode? At issue is whether the birds could use this new general knowledge about degradation profiles to alter their recovery preference, even though this new information was not available at the time when the cache memory was encoded.

To address this issue, following the test of recovery after different retention intervals we gave our jays interleaved caching and recovery trials using the design illustrated in Fig. 9.3




Fig. 9.3 The experimental design of the interleaved trials experiment for the reversed and consistent groups. P: peanuts; C: fresh crickets; dC: degraded crickets; ?: on test, all caches from both sides of the tray are removed to establish where the birds will search in the absence of any cues directly emanating from the food.

(Clayton et al. 2003). Although the design looks complex, the basic idea was to give the jays in the reverse condition information that the crickets did in fact decay during the retention interval. So these birds were allowed to cache crickets and peanuts in each of three different trays on successive days. For example, a jay might cache the two foods in tray 1 on Monday (day 1), and then cache more crickets and peanuts in a different tray 2 on Tuesday (day 2) and in a third tray 3 on Wednesday (day 3). Remember that none of the birds had been given any information about whether crickets are fresh or degraded after 3 days because they had been trained only with retention intervals of 1 day and 4 days. The birds were then provided for the first time with direct information about the fate of their cricket caches after a 3-day retention interval so that this is information was available only after they had finished caching in all three trays. Therefore, on Thursday (day 4) the reverse group recovered peanuts and crickets from Monday's tray 1 and, against their expectation, discovered the crickets had in fact perished after the 3-day retention interval. This new information was reinforced by recovery from Tuesday's caching tray 2 on Friday (day 5) before these jays were given a critical test in which they searched in Wednesday's tray 3 on the Saturday (day 6). Of course, being a test trial no food was present at recovery. The recovery preferences of the birds in this reverse condition on this test trial was contrasted with that of another set of birds from the degrade group. The jays in this consistent condition had their expectation that the cricket would be fresh after 3 days confirmed. Searching in tray 1 on Thursday (day 4) and tray 2 on Friday (day 5) yielded fresh and palatable crickets.

Predicting the search preference for the consistent group is straightforward. The experience of recovering palatable crickets from trays 1 and 2 should have confirmed the jays' generalized belief(s) about the decay profile for crickets and therefore these birds should have searched for crickets in the left side of tray 3 on the test recovery period on day 6 (Fig. 9.3). And this is what all four of these jays did. For the reversed condition, however, the experience of recovering perished crickets from trays 1 and 2 on days 4 and 5, respectively, should have disconfirmed the jays' beliefs that crickets were fresh after 3 days and replaced this general knowledge belief with one representing the crickets as degraded after this retention interval. Consequently, when the episodic-like memory of caching in tray 3 on day 3 was retrieved by the presentation of this tray on day 6, the interaction of this memory with the birds' general knowledge through the practical inference processes should have derived an intention to search for peanuts. A further interaction with the episodic-like memory of where peanuts were cached in tray 3 should have lead to an intention to search in the left side of tray.

Indeed, all four birds in the reversed condition directed their first search during the test on day 6 to the peanut side of tray 3. Importantly, this switch in preference did not reflect a general change in their beliefs about whether or not to search for crickets because when the birds were tested with a 1-day retention interval they reverted to their prior preference for crickets. We therefore interpret this reversal of the search preference as evidence that our jays can integrate information about the caching episode with new information presented during the retention interval in a rational manner.

9.4 The rationality of caching strategies

As we have emphasized, the evaluation of the psychological rationality of animal action is a matter of bringing converging evidence to bear of the issue because intentionality is not necessarily manifest in behaviour. It is often claimed that the demands of social interactions, and especially competition from conspecifics, is a major factor in the evolution of cognition (Humphrey 1976; Jolly 1966). In this respect, it is notable that for some species, food caching and recovery are activities that occur within a social context, not least because caches are susceptible to pilfering by other individuals (Van der Wall 1990). Many species are known to pilfer the caches of other birds, and several species of corvid, including our scrub jays, use observational memory to locate the caches of their competitors and pilfer them when the food-storer has left the scene (see Clayton et al. 2001). But these food-caching corvids also engage in a number of strategies that serve to reduce the probability that their caches will be stolen by competitors. For example ravens will delay caching if other ravens are in the vicinity, waiting until the potential pilferers are distracted or have left the scene (Heinrich and Pepper 1998; Bugnyar and Kotrschal 2002). In other instances they will preferentially store food behind obstacles so that other ravens cannot see where the caches are being made (Bugnyar and Kotrschal 2002). But when food is readily available and many other individuals are present it

might be better to cache while you have the opportunity to do so, in the hope that your competitors are distracted! There are a number of reports of corvids, including ravens and jays, caching rapidly when other individuals are present, only to return alone when the competitors have left the scene, moving or recaching the caches they had hidden in the presence of conspecifics to new cache sites (e.g. Bugnyar and Kotrschal 2002; Emery and Clayton 2001; Goodwin 1956; Heinrich 1999).

As noted earlier, the apparent complexity of a particular behaviour, in this case caching, does not speak to the issue of whether or not the behaviour is rational. So is there any evidence that our scrub jays can make rational decisions about the future consequences of caching and stealing? Consider the apparent observation of birds moving the food they had hidden in the presence of other individuals, and recaching the food items in new places when those observers were no longer present. Of course in the wild, one might explain the presence or absence of another bird as purely coincidental to the caching and recaching events. So first, one needs to test whether it is the presence of an observer at caching, and absence of one at recovery, that elicits the food-cacher's recaching behaviour. This is what we set out do in the first test (Emery and Clayton 2001) illustrated in Fig. 9.4.

We allowed our scrub jays to cache food under two conditions, namely in one uniquely-configured tray when watched by another jay (the observer), or in another




Fig. 9.4 The experimental design of the first social context of caching experiment, where storer jays cached worms in unique trays either when observed by another jay (observer) or in private (with the view of the observer occluded. After a 3-hour period, the storers recovered their caches in private, independent of the social context during caching. An additional, second unique tray provided opportunities for re-caching.

uniquely-configured tray when a cover was placed over the other bird's cage so that the storer could hide its food in private. Subsequently, the storer was presented with the tray in which it had cached (old tray) as well as a new tray (also with a unique arrangement of toy blocks) in which it could recache, and allowed to recover and recache the food in private so that at the time of recovery the conditions were identical. In other words, the only difference between the two treatments was in whether the jay had been observed during the caching phase or had been allowed to hide its food in private. If recaching really depends on the presence of an observer at the time of caching, then we should only expect to see recaching if the bird had been observed, and not when it had been allowed to cache in private. And so it was in this experiment. The jays continued to cache food, irrespective of whether another bird was watching them or not. The key result was that birds recached significantly more items during recovery when they had been observed during caching. Moreover, they specifically placed these recovered caches in new cache sites the observer had not seen, but only when they had been observed caching. When they had cached in private, they were just as likely to place the few recaches that they did make in original cache sites as in the new ones. Because the two conditions were identical at the time of recovery, the difference between treatments meant that the birds had to remember whether or not they had been watched during the previous caching condition in order to know when to recache during recovery, and whether to recache in new sites.

Before considering an intentional account of recaching, we attempted to rule out some more obvious mechanistic explanations. For example generalized recaching may simply be triggered automatically by the memory of being in the presence of another bird recently. Consequently in the next study, we investigated whether recaching is mediated by an episodic-like memory of the specific caching episode during which the jay was observed. To do so, we used an interleaved trials design similar to that described for the reversed and consistent conditions of the cricket degradation study, except that this time we gave the birds 'observed' and 'in private' caching trials (Fig. 9.5). Thus some birds cached in the observed tray while watched by a conspecific and then cached in the other, in-private tray immediately afterwards. The remaining jays received these two caching episodes in the reverse order. Subsequently the storers were allowed to recover from both trays in the presence of a novel tray in which they could recache. The birds mainly recached from the observed tray, moving them to new sites. By contrast, the birds showed little tendency to recache items from the in-private tray, nor did they discriminate between old and new sites for those items that they did recache. So this result suggests that jays remember not only whether or not they were being observed, but can also recall the specific tray in which they cached while being watched, rather than relying on a default mechanism to recache if they have been observed recently.

The obvious intentional account of recaching assumes that the birds have general beliefs that being observed during caching causes the loss of the cached food items and

212 RATIONAL ANIMALS?




Fig. 9.5 The experimental design of the interleaved trials experiment on the social context of caching. Storer jays either cached in one unique tray when observed (observed tray) then 10 min later in a second unique tray in private (in private tray) or *vice versa*. After a 3-hour period, the storers were allowed to recover their caches from the observed and in private trays outside of the view of the observer (in private). An additional, third unique tray was provided for re-caching.

that this loss can be prevented by recaching. When taken in conjunction with a low desire for pilfered caches relative to intact caches, the process of practical inference would rationalize recaching with respect to the content of these beliefs and desires. However, this account immediately raises the issue of where these beliefs come from; unless one is prepared to countenance the idea of innate beliefs with intentional content (as opposed to innate associative structures), the answer must lie with the past experience of the storer. Therefore, in a final study, we compared recaching by the jays described above, all of whom had experience of caching and of observing and subsequently stealing other birds' caches, with two other groups of jays (Emery and Clayton 2001). The observer group had experience of watching other birds cache, but had never been given the opportunity to steal those caches. By contrast, the pilferer group had experience of stealing other birds' caches, but were not tested explicitly for their ability to remember the location of those caches within an experimental context. (We cannot say that these jays had no experience of observing other birds caching outside of the experiments.) We found that experience did matter. Indeed, jays that had prior experience of pilfering another bird's caches subsequently recached food in new cache sites during recovery trials, but only when they had been observed caching. Jays without this

pilfering experience did not, even though they had observed other jays caching. These results therefore suggest that jays relate information about their previous experience as a thief to the possibility of future stealing by another bird, and modify their caching strategy accordingly.

We draw two conclusions from this study. The first is that recaching is psychologically rational in that, of the explanations currently available, only an intentional account explains the differences in recaching behaviour between those with and without experience of having been a pilferer in the past. Second, it is important to note that recaching appears to be based on mental attribution or 'mind reading' (see this volume, Part V).³ The inference that jays with prior pilfering experience appear to make in this situation is that a conspecific with similar prior experience would share the same beliefs as they have, namely that caches can be stolen. It is this inference that both rationalizes and causes recaching. This conclusion, of course, goes beyond the issue of the simple behavioural rationality and clearly requires further examination.

9.5 **Summary and conclusions**

In this chapter we have reviewed some of our studies of cache recovery and recaching by western scrub jays with respect to the issue of whether these behaviours are psychologically rational. We have assessed psychological rationality by investigating whether the jays' behaviour warrants an intentional belief-desire explanation. With respect to desires, motivational manipulations that should have yielded differential desires for two types of cached foods caused the birds to search preferentially for the most desired food in a task in which a simple associative account predicted the opposite preference. Moreover, the response to these motivational manipulations provided a procedure for investigating the content of the birds' desires.

Our analysis of the beliefs underlying cache recovery was set within the framework of a declarative memory system that distinguished between general beliefs or knowledge and episodic memories. According to this analysis, the information about a specific caching event is represented in a what-where-when memory with episodic-like properties, whereas general knowledge about the relative desirability of difference food types and the way in which desirability changes with the time in the cache, is represented by semantic-like memories. Searching for caches at recovery is then generated by the interaction of these two types of memory through processes of practical inference.

³ Editors' note: It is interesting to assess this evidence for mental attribution in light of the methodological constraints urged by Povinelli and Vonk, this volume. Does it satisfy their concerns? It can be argued that this is some of the best evidence to date precisely because of the levels of control: (a) recovery conditions are always conducted in private and therefore any difference between treatments must result from differences at the time of caching, hours earlier and (b) only some of the jays—and those with specific prior experience of being pilferers—engage in this behaviour. The editors are grateful to the authors of this chapter for discussion of this point.

It is the nature of this interaction that gives cache recovery its rational character and endows this behaviour with a degree of flexibility. This flexibility was demonstrated in procedures in which we gave the jays new general information during the retention interval that was not available to the birds at the time of caching. The reversed search preference at recovery demonstrates that they were capable of integrating this new information with their memory of the specific caching episode in a way predicted by the rational–declarative model.

Further evidence for mnemonic integration comes from studies demonstrating that the jays can integrate information from their own experience of pilfering other birds' caches with memories of caching episodes in which they themselves were observed caching food. This integration produced recaching of food items into new sites, a rational defensive action on the part of the storer. Although none of these examples of the flexibility of caching and recovery behaviour is decisive on its own, taken together they provide strong, converging evidence for the intentional and rational control of behaviour especially when contrasted with the failure of standard associative accounts.

Many, of course, have claimed that mammals, and especially primates (see Call, Chapter 10; Tomasello and Call, Chapter 17; Boysen, Chapter 22, this volume) and cetaceans (see Herman, Chapter 20; Tschudin, Chapter 19, this volume) are capable of rational cognition, and indeed one of us (AD) has argued that even the goal-directed behaviour of rodents is mediated by intentional processes (Heyes and Dickinson 1990; Dickinson and Balleine 2000). Our research on caching and recovery by western scrub jays contributes to the growing body of evidence that behavioural markers of rational cognition are as compelling for at least some birds (see Pepperberg, Chapter 21; Papineau and Heyes, Chapter 8; Kacelnik, Chapter 2, this volume) as they are for mammals in spite of their radically divergent neurobiology and evolutionary trajectory (see Emery and Clayton 2004, for a detailed discussion).

Acknowledgments

This research was supported by grants from National Institute of Health (NS35465 and MH62602), and the Grindley Trust to Nicola Clayton, a BBSRC Grant (S16565) to Nicola Clayton and Anthony Dickinson, and conducted within the UK Medical Research Council Co-operative Grant (G9805862). Nathan Emery was supported by a Royal Society University Research Fellowship.

References

- Bugnyar, T., and Kotrschal, K. (2002). Observational learning and the raiding of food caches in ravens, *Corus corvax*: Is it 'tactical' deception? *Animal Behaviour*, **64**: 185–195.
- Clayton, N. S., and Dickinson, A. (1999). Memory for the contents of caches by scrub jays (*Aphelocoma coerulescens*). *Journal of Experimental Psychology: Animal Behavior Processes*, **25**: 82–91.
- Clayton, N. S., and Griffiths, D. P. (2002). Testing episodic-like memory in animals. In: L. Squire and D. Schacter, eds. *The Neuropsychology of Memory*, 3rd edn, pp. 492–507. New York: Guilford Publications.

- Clayton, N. S., Griffiths, D. P., and Dickinson, A. (2000). Declarative and episodic-like memory in animals: Personal musing of a scrub-jay. In: C. Heyes and L. Huber, eds. *The Evolution of Cognition*. Cambridge, MA: MIT Press.
- Clayton, N. S., Griffiths, D. P., Emery, N. J., and Dickinson, A. (2001). Elements of episodic-like memory in animals. *Philosophical Transactions of the Royal Society (London)*, B, **356**: 1483–1491.
- Clayton, N. S., Yu, K. S., and Dickinson, A. (2001). Scrub-jays (Aphelocoma coerulescens) form integrated memories of the multiple features of caching episodes. Journal of Experimental Psychology: Animal Behavior Processes, 27: 17–29.
- Clayton, N. S., Yu, K. S., and Dickinson, A. (2003). Interacting cache memories: evidence of flexible memory use by scrub jays. *Journal of Experimental Psychology: Animal Behavior Processes*, **29**: 14–22.
- Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society (London)*, B, **308**: 67–78.
- Dickinson, A. (1989). Expectancy theory in animal conditioning. In: S. B. Klein and R. R. Mowrer, eds. *Contemporary Learning Theories: Pavlovian Conditioning and the Status of Traditional Learning Theories*, pp. 279–308. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Dickinson, A., and Balleine, B. (1994). Motivational control of goal-directed action. Animal Learning and Behavior, 22: 1–18.
- Dickinson, A., and Balleine, B. W. (2000). Causal cognition and goal-directed action. In: C. M. Heyes and L. Huber, eds. *The Evolution of Cognition*, pp. 185–204. Cambridge, MA: MIT Press.
- Dickinson, A., and Balleine, B. (2002). The role of learning in the operation of motivational systems. In: H. Pashler and R. Gallistel, eds. *Stevens' handbook of Experimental Psychology, vol. 3, Learning, motivation, and emotion*, 3rd edn, pp. 497–533. New York: John Wiley and Sons.
- Emery, N. J., and Clayton, N. S. (2001). Effects of experience and social context on prospective caching strategies in scrub jays. *Nature*, **414**: 443–446. [See also *Nature*, **447**: 349.]
- Emery, N. J., and Clayton, N. S. (2004). Comparing the complex cognition of birds and primates. In: L. J. Rogers and G. S. Kaplan, eds. *Comparative Vertebrate Cognition: Are Primates Superior to Non-Primates?*, pp. 3–55. The Hague: Kluwer Academic Press.
- Fodor, J. A. (1977). The Language of Thought. Hassocks, UK: Harvester Press.
- Goodwin, D. (1956). Further observations on the behaviour of the jay. Ibis, 98: 186-219.
- Griffiths, D., Dickinson, A., and Clayton, N. S. (1999). Episodic memory: what can animals remember about their past? *Trends in Cognitive Science*, **3**: 74–80.
- Heinrich, B. (1999). Mind of the Raven. New York: Harper Collins.
- Heinrich, B. (2000). Testing insight in ravens. In: C. M. Heyes and L. Huber, eds. *The Evolution of Cognition*, pp. 289–309. Cambridge, MA: MIT Press.
- Heinrich, B., and Pepper, J. (1998). Influence of competitors on caching behaviour in common ravens, *Corvus corax*. *Animal Behaviour*, **56**: 1083–1090.
- **Heyes, C., and Dickinson, A.** (1990). The intentionality of animal action. *Mind and Language*, 5: 87–104.
- Humphrey, N. K. (1976). The social function of intellect. In: P. P. G. Bateson and R. A. Hinde, eds. *Growing Points in Ethology*, pp. 303–317. Cambridge: Cambridge University Press.
- Jolly, A. (1966). Lemur social behaviour and primate intelligence. Science, 153: 501–506.
- Nagel, T. (1974). What is it like to be a bat? Philosophical Review, 83: 435-450.
- Rumelhart, D. E., and McClelland, J. L. (1986). Parallel Distributed Processing. Cambridge, MA: MIT Press.
- Thorndike, E. L. (1911). Animal Intelligence: Experimental Studies. New York: Macmillan.
- **Tulving**, E. (1972). Episodic and semantic memory. In: E. Tulving and W. Donaldson, eds. *Organization of Memory*, pp. 381–403. New York: Academic Press.

216 RATIONAL ANIMALS?

Tulving, E. (1983). Elements of Episodic Memory. New York: Oxford University Press.

Tulving, E. (2001). Episodic memory and common sense: how far apart? *Philosophical Transactions of the Royal Society (London)*, **B356**: 1505–1515.

Tulving, E., and Markowitch, H. J. (1998). Episodic and declarative memory: role of the hippocampus. *Hippocampus*, 8: 198–204.

Van der Wall, S. B. (1990). Food Hoarding in Animals. Chicago: University of Chicago Press.

Winograd, T. (1975). Frame representations and the declarative-procedural controversy.

In: D. Bobrow and A. Collins, eds. *Representation and Understanding: Studies in Cognitive Science*, pp. 185–210. San Diego, CA: Academic Press.