11 Selective Imitation in Child and Chimpanzee: A Window on the Construal of Others' Actions

Andrew Whiten, Victoria Horner, and Sarah Marshall-Pescini

11.1 Introduction

Suggesting that the behavioral repertoire of human children is shaped by the particular culture they are reared in, and that this occurs in large measure through various kinds of broadly imitative, "tradition-copying" processes, is hardly likely to provoke much debate. That the same might be true of our closest relatives, chimpanzees, has been much more controversial.

The possibility of some such fundamental similarities between chimpanzees and ourselves became a working assumption following the first substantial field studies of chimpanzees (Goodall, 1973), reinforced by impressive-looking experimental evidence for a wide-ranging imitative capacity (Hayes & Hayes, 1952). However, these conclusions were later challenged by authors who argued that the evidence was not in fact adequate to demonstrate that chimpanzees were either truly cultural, or imitative (Tomasello, 1990; Galef, 1992; Heyes, 1993), a view reinforced by new experimental studies reporting a lack of imitative copying of tool use (Tomasello et al., 1987). These and other critiques have been followed by a decade of effort in which the chimpanzee data have become more rigorously obtained and substantial. In particular, the charting of putative chimpanzee cultural variation has become more systematic (Whiten et al., 1999, 2001) and several careful experimental studies have followed the provocative findings of Tomasello et al. (1987); reviewing the field, Whiten et al. (2004) were able to cite thirty-one experimental studies completed in the intervening period. A measure of the impact of this work is that the authors of some of the sternest earlier critiques have been prepared to revise their conclusions. Thus experimentalist Tomasello joined with field ethologist Boesch to agree that "a comparison of chimpanzee and human cultures shows many deep similarities" (Boesch & Tomasello, 1998, p. 591), and

Heyes (2001a, p. 253) concluded that recent experimental studies had demonstrated chimpanzee imitation.

So it seems that "apes ape," after all, and the idea of chimpanzee cultures has become widely accepted. To a casual observer it may seem that the field has thus merely come full circle, to a picture of chimpanzee–human similarities that many long assumed existed. Indeed, broadly defined, both imitation (in Thorndike's (1898) sense of "learning to do an act from seeing it done") and "culture" (in the sense of "behavioral traditions") appear to be quite widespread among birds and mammals (Zentall, 2001; de Waal, 2001a), so that claiming the existence of both in our closest relative could be dismissed as not so very startling.

This, however, would be to neglect the fact that our current knowledge is built on much firmer scientific foundations than before, and that we have gained a significantly richer picture of the complex nature of both traditions and imitation in chimpanzees. Indeed in chimpanzees as in humans, there is more to "culture" than the mere existence of traditions (Whiten, 2000, 2001; Whiten et al., 2003). In a similar way, there can be more to imitation than implied by Thorndike's basic definition, and in this chapter we address some of these complexities in the imitation of human and nonhuman apes. Our central point is that even though a substantial cultural repertoire may be acquired by imitative copying, neither children nor chimpanzees copy all they see others around them doing. Indeed, a moment's reflection confirms they could not; if they did, they would appear to be little marionettes perpetually mimicking whoever was in view. So the question naturally arises, What determines what is copied, and what is not? This is a question little addressed by research as yet, but it is surely an important one. Thus in relation to the mechanism, or process, of imitation that is addressed by so many chapters in these two volumes, we are led to examine most fundamentally *what the imitative process selectively imitates.*[1] This in turn raises an important functional question of what purposes any such selectivity serves. Finally, this issue is important in relation to the topic of cultural transmission mentioned earlier and covered elsewhere in this volume and in volume 2, for selectivity in imitation clearly has significant potential to shape cultural stability and change.

We have been led to the question of selectivity principally because in so many of our studies, when we find evidence of imitation in apes, it is typically *partial*; some parts of what a model does are copied, others are not.

1. For relevant discussion, see chs. 16 by Harris and Want and 17 by Greenberg in vol. 2. ED.

This is particularly the case in chimpanzees compared with children. One interpretation of such findings, particularly likely to recommend itself to those skeptical of imitation by apes, is that this merely reflects a relatively poor imitative ability, and this is one theoretical possibility we address here. However, we also find striking biases in what children do or do not imitate and some of these, as in the chimpanzees, are suggestive of active and functional selectivity. Clearly, we need a method of distinguishing any such true and interesting selectivity from patchy copying caused merely by limited imitative competence.

Our answer has been to design experiments in which we systematically *predict* what subjects will or will not selectively copy in two contrasting conditions, *if* they are operating on adaptive or functional grounds, as opposed to being merely limited in their imitative capacity. If we can successfully predict the conditions under which particular aspects of a model's actions will be reproduced, we are getting onto firm ground in concluding that true selectivity is operating. In addition, to the extent our predictions are derived from a systematic theoretical framework, we will have made some real progress in understanding *why* the selectivity takes the form it does. We therefore explain our framework next.

11.1.1 A Framework for Analyzing Selective Imitation

Our framework for these studies draws on theoretical discussions in the literature of the past 15 years or so (Whiten et al., 2004) and is illustrated in figure 11.1. As indicated, we envisage the processes involved in copying from another's actions as constituting a continuum ranging between imitation and emulation. The core idea we wish to capture is that at the emulation pole of the continuum, the observer may be actively ignoring—"selecting out"—various aspects of the actions of the model, which at the imitative pole would instead be "selected in." We can envisage various contexts in which such a flexible, selective strategy could be adaptive, as indicated in figure 11.1. For example, some actions may be seen as accidental rather than intended, or as not causally linked to an outcome of interest, and thus not copied. Or it might be that even though certain acts *are* perceived as intended and/or causally necessary, the observer has at his or her disposal alternative behavioral techniques for gaining the results of interest, which are preferred over those performed by the model (and of course even if components are selected out because they are seen as *not* intended or *not* causal, emulation requires that the observer be able to come up with actions that can be substituted for those used by the model).

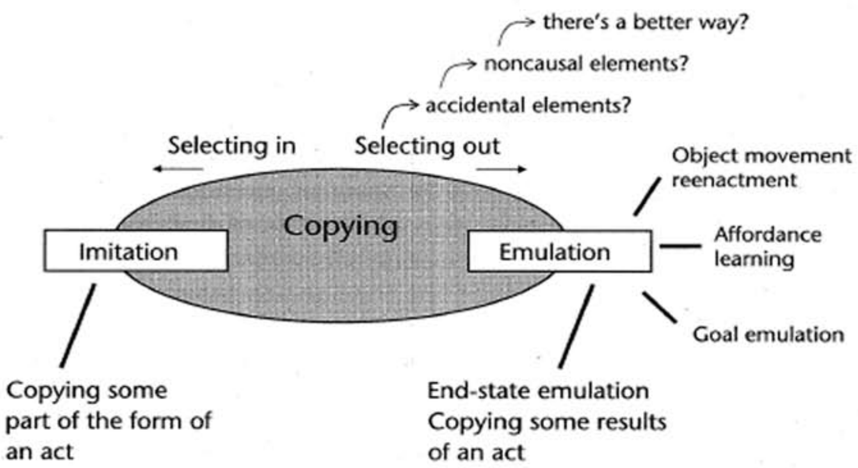

Figure 11.1
Framework for analyzing selective imitation. Copying is conceptualized as a continuum. At one pole, emulation may involve copying only the results the model achieved, or their goal, or the way objects are moved; at the other pole, imitation involves a more complete copy of the form of the actions of the model. Thus in emulation there is a relative selecting out of certain parts of the model that at the imitative pole are selected in. Three possible bases for selecting out elements of the model are suggested: they are perceived as accidental rather than intended; they are perceived as noncausal rather than causing a desirable outcome; or the observer simply knows a better way to achieve such an outcome (see the text and Whiten et al., 2004, for further discussion and references).

The hypothesis we wish to explore goes beyond a common preoccupation in the recent literature about whether any particular species is an imitator or an emulator, to propose that instead there may be adaptive flexibility, so that a learner might switch between the imitation and emulation poles of the copying continuum according to context. In this chapter we focus on the opportunities to select in or select out information on the causal relevance of actions the model performs. Our previous research suggests that imitative copying is favored in tasks where such causal connections are relatively opaque to the learner; under these constraints, it may be adaptive to imitate rather comprehensively all the actions taken by the model. By contrast, where causal connections are more transparent to the observer, emulation becomes more feasible. Emulation may be preferred for reasons such as those suggested earlier, and may be more efficient

than what can be achieved through the higher-fidelity copying that distinguishes the imitation end of the continuum.

We have turned these ideas into experiments that we describe in more detail in the following sections. The basic strategy is to create two different forms of the same task, one of which we predict will generate relatively full imitation, and the other a more restricted copy limited to end results of the sequence of actions. Such an approach not only probes the mechanisms of observational learning, but can also give us, as we indicate in our subtitle, a window on the construal of others' actions. Thus in the present case, they may tell us something of the extent to which chimpanzees' and children's perceptions of others actively assess the causal significance of the actions performed and use this assessment to shape the kind of social learning process activated. This issue is given added topicality because of a suite of experimental studies by Povinelli (2000) that led him to conclude that chimpanzees actually understand rather little about causality.

11.2 Imitation in Chimpanzees: Ill Formed or Selective?

Before describing the experiments designed explicitly to test the hypothesis outlined earlier, repeated findings in earlier work need to be reviewed briefly. The basic finding of interest in our earlier studies on evidence of imitation, whether in chimpanzees or children, is that the copying tends to be partial. Overall, especially in chimpanzees, we have tended to find patchy fidelity to the actions of the model. As noted earlier, it is often inherently difficult to discern whether this reflects an imperfect imitative ability or is instead demonstrating the kind of adaptive selectivity described in the foregoing section. This is why we have performed the prospective experiments described here, which were specifically designed to differentiate these possibilities. Nevertheless, it is important to review certain earlier findings precisely because they do form the background and platform for the new experiments; moreover, they repeatedly present puzzles of interpretation that may be solved at least in part by the theory and experimental approaches described in this chapter.

11.2.1 Copying Some Components of a Task, but Not Others

In three studies of the observational learning of chimpanzees we have used an "artificial fruit" (figure 11.2), designed as an experimental analogue of foraging problems encountered in the wild (Whiten, 1998, 2002a; Whiten et al., 1996). As with certain real fruits, the animal is faced with a task

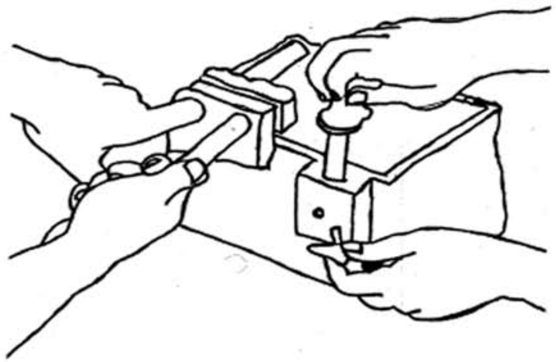

Figure 11.2
"Pin-apple" artificial fruit for experimental investigation of imitation. The actions illustrated are poking through the outer bolt, pulling and twisting out the inner bolt, removing the pin, and pulling up the handle. Once all defenses are removed in these ways or others, the lid can be opened to gain access to an edible core.

that involves removal of the food's defenses (e.g., shelling, breaking off) to allow access to an edible core. In addition, the fruit was designed to implement a two-action approach (B. Dawson & Foss, 1965) to assessing imitation. Each defense can be removed in one of two ways, only one of which is seen by each subject before they are given the chance to try to open the fruit themselves. The extent to which the two groups of subjects subsequently perform differently, and in particular the extent to which their actions match those of the model variant they watched earlier, gives us a robust measure of imitative learning. Differences cannot be due to stimulus enhancement, since both of the two variant techniques are applied to the same part of the fruit.

We have consistently found that chimpanzees performed imitative copying of certain parts of this task and not others (Whiten, 1998, 2002a; Whiten et al., 1996). The two alternatives of *poking* the "bolts" out of the way, versus *pulling and twisting* them out (see figure 11.2), elicited matching behaviors in chimpanzees. These behaviors could be identified by videotape coders blind to which variant the subjects had actually seen. Such results were similar to those for preschool children tested in similar ways (in these experiments, adult human models who were familiar to the subjects were used in all cases). By contrast, chimpanzees that had witnessed a model either pulling out the "handle" (as in figure 11.2) or twisting it out of the way, did not show significantly different approaches when they had

their turn at the task; their tendency was efficiently to pull out the bolt, whichever method they had seen modeled. From the perspective of figure 11.1, the chimpanzees' actions were relatively *imitative* with respect to the bolts, but relatively *emulative* with respect to the handle; for the latter they appeared to have learned something about a useful end result (disabling the handle defense), but applied their own technique to achieving this. However, children continued to imitate this part of the task, as they had with the bolts.

The basic contrast between these two sets of subjects is thus that the children exhibited imitation of a quite high fidelity throughout, whereas the chimpanzees imitated only partially—on the bolts but not the handle. One interpretation of these results is that chimpanzees are simply less proficient imitators. However, we found that the chimpanzees were in fact faster at opening the fruit, apparently because they did not copy all they had seen in the relatively painstaking way characteristic of the children. Instead, by applying their own preferred method of disabling the handle defense, they turned in a more efficient performance than the children. This argues for seeing the chimpanzees' approach as not necessarily expressing an inability, but instead a possible case of adaptive, selective imitation.

This appears to leave us with the question of why the children were not acting in such an adaptive way, a puzzle we postpone for discussion after some other relevant results have been outlined. For the chimpanzees, the question is why they copied one component and not the other. Our guess is that from the chimpanzees' perspective, the two components differed in transparency. As chimpanzees acted on the handle, even if they were turning it, they would tend to discover it could be pulled out, and they preferred to do so subsequently (unlike the more faithfully copying children). By contrast, if the chimpanzees started to poke out the bolts when they had seen the model do this, it appeared to be not so transparent to them that the bolts could in fact be pulled out instead.

This interpretation of the chimpanzees' approach as adaptive, selective imitation of some components but not others is, of course, post hoc. It can be tested only by further experiments of the kind we report further here, but experiments informed by the kinds of results we have just described.

11.2.2 Differential Copying of the Sequential Structure of a Technique versus Its Components

Extending the concept of the two-action methodology (Whiten, 2002a), the artificial fruit described earlier was designed so the steps to opening it could be performed in different sequential orders (figure 11.2). In an earlier study, two chimpanzees saw each of the two bolts removed first; then the pin was taken out, the handle was disabled, and the lid opened to reveal the edible part inside. Another pair of subjects saw the alternative sequence of pin-handle-outer bolt-inner bolt. Within each pair, one subject saw one technique for each component (e.g., poke the bolts) and the other subject saw the alternative (e.g., pull and twist the bolts) (Whiten, 1998).

By the third and final trial, there was a statistically significant match to the particular sequential structure each subject had witnessed. In three cases the match was perfect and in the fourth the only difference was the relatively trivial one that the bolts were taken out in a different order. By contrast, these subjects varied in the extent to which they matched details of the ways in which the component acts had been performed. Two showed quite extensive matching of techniques applied to the bolts; one replicated the spinning action it had seen on the pin and one subject did neither of these. Thus, in these results we may be seeing selective imitation (to varying extents in different individuals) of the overall sequential structure, at the expense of details of the technique for each component. This might constitute adaptive selectivity to the extent that in many complex natural tasks it may be more critical to learn the overall structure than the individual actions (R. Byrne & Russon, 1998), at least when the animal is competent enough to generate its own local solutions to the latter.

However, if we take a less static view and consider the development, or learning, of a skill, the selectivity may be more complex than this. As noted earlier, the subjects came to match the sequential structure only after several trials (whether this reflects repeated opportunities to observe the model, or repeated opportunities to perform the task, or an interaction between these, we do not yet know). By contrast, there was evidence of copying component techniques (twisting bolts and spinning the pin, neither of which are in fact physically necessary) from the outset. This raises the possibility that there is differential selectivity during the social learning of a complex sequence of actions, with selectivity initially focusing on the shape of the component acts, then later shifting to higher-level sequential patterning (Whiten et al., 2004). This might vary among species. Testing gorillas, we found no evidence of sequence learning in equivalent tests, yet there was evidence that they copied aspects of the components, such as spinning the pin versus pulling it straight out, the two variants demonstrated by models (Stoinksi et al., 2001).

As in the previous study, however, we must be cautious about these interpretations. It remains possible that the apparent "selectivity" observed in this study in fact reflects merely individual differences in imitative competence, rather than active selectivity. We have here important background findings that require the new experiments we report in the following discussion, and are no substitute for them.

11.2.3 Differential Copying of Hierarchical versus Sequential Structure

Imitation might occur at the level of the hierarchical structure that characterizes skills such as foraging techniques among apes (R. Byrne, 1994 and vol. 1, ch. 9). Although the copying of sequential structure described earlier may be an instance of this, it is possible, at least in principle, that it reflects only the copying of a linear, sequential chain of actions, A-B-C, etc. To actually test for the copying of hierarchical structure, we need materials specifically designed for this purpose (Whiten, 2002a,b).

Working with preschool children, we have used a "key-way" "fruit" that requires sixteen component actions to allow a lid to be finally removed to obtain a reward (Whiten, 2002a,b; see also figure 11.3). The sixteen actions can be performed by models in one of two different, hierarchically organized ways. Imagining the sixteen actions as a 4×4 array, one can characterize the two different approaches as either "column-wise" or "row-wise." In the column-wise approach, a handle is selected and stabbed into a shaped tablet to make a "key"; this is then inserted into an appropriate lock (figure 11.3), pushing a restraining rod out of the way, which is in turn removed. This process is repeated a further three times with new handles and tablets, until all the restraining rods are removed and the lid can be lifted off. In the alternative row-wise approach, all the handles are picked up and stabbed into the tablets in turn; all the keys are then inserted, and finally all the rods are removed and the lid lifted off. It was found that preschool children were really quite faithful in copying whichever of the two overall approaches they had seen, which provides the first direct evidence of the imitation of hierarchical structure in any species.

However, each of the two sets of subjects had been further subdivided. Half of each group saw the components of the task performed left-to-right and half saw them done right-to-left. Thus, in addition to testing for imitation of overall hierarchical structure, we could examine whether the particular, chainlike sequence of components modeled was copied. We found that in these young children it was not. There was a stark contrast between copying of the hierarchical structure itself and ignoring of the particular sequential chain of component actions witnessed.

Presenting this task to chimpanzees is problematic because its small parts are so easily stolen and chewed! One of us (SM-P) therefore designed

Figure 11.3
"Key-way" fruit for identifying imitation of hierarchical structure. A lid is held on an underlying box by four horizontal rods that can be seen through the transparent lid. One way in which a model opens the fruit is by taking one of the handles from the central cup, stabbing it into a tablet held in a recess at the back of the lid and thus making a key, pushing this key into a correspondingly shaped slot at the front (illustrated here) to eject one of the rods, pulling out the rod thus protruding from the back of the box, and finally removing the key. A similar "column" of actions is performed for each of the shapes indicated (– X T L). In the alternative "row" method, all the handles are first stabbed into all the tablets; then all the keys are placed in the appropriate locks, all the rods are pushed through and taken out, and finally all the keys are removed. The two approaches thus include the same set of sixteen actions, organized into two alternative hierarchical structures ("columns" versus "rows").

another task to test for copying hierarchical structure, built on the same principles as the key-way, but somewhat simpler in that it incorporated only a 3 × 4 array of component actions (figure 11.4). The column-wise approach to this task involves removing a bolt of the kind used in the pin-apple described earlier, folding back the hinged door that had been held shut by the bolt, flicking aside a sliding door, and turning a knurled knob underneath it several times. Once all three columns are completed, a drawer opens, offering a reward. The row-wise approach, of course, involves performing all bolt removals, then all door folds, and so on to overall completion.

Juvenile chimpanzees have shown a significant tendency to follow whichever of the two approaches they have seen (Marshall-Pescini & Whiten, unpublished), which to our knowledge is the first quantitative

a

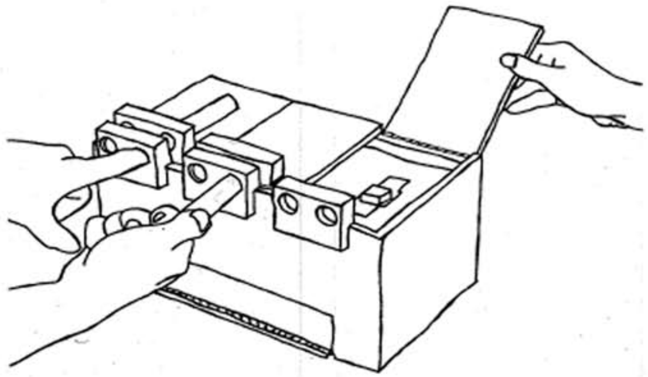

b

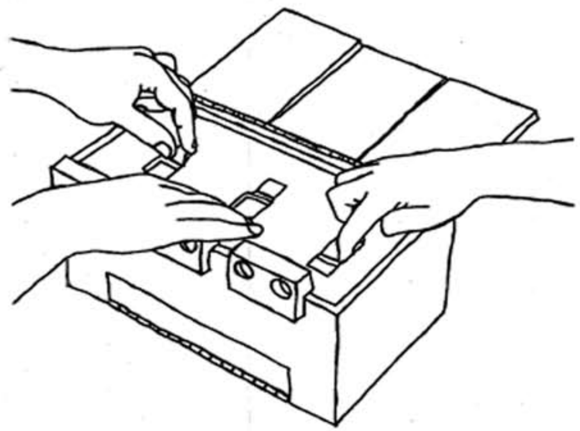

Figure 11.4
Task used for identifying imitation of hierarchical structure in chimpanzees. (*a*) The actions illustrated are poking a bolt through (left), pulling and twisting it out (center), and opening the lid once the bolt is removed (right). (*b*) Additional actions illustrated are flicking the sliding door away (left), flicking it in the opposite direction (middle), and inserting a finger to turn the knurled knob underneath (right). Once all twelve operations are completed, the door at the base drops open to reveal food. See the text for a further explanation of alternative hierarchical approaches to the task.

evidence for copying of hierarchical structure in a nonhuman primate. It thus complements the similar finding for children on the key-way fruit described earlier. There is also a similarity with the earlier study of sequential imitation (Whiten, 1998) in that a significant match emerged only after several trials. However, as in the key-way study, the chimpanzees were exposed to different chains of actions (working left-to-right versus right-to-left) within each of the hierarchical approaches. And like the children, the chimpanzees did *not* copy this specific chain of actions, even though they picked up the higher-level hierarchical structure they saw modeled.

The results of the child and chimpanzee versions of this task thus show a remarkable convergence. In both cases we appear to have selective imitation of hierarchical structure, ignoring sequential chaining of component actions. As before, we must of course consider the possibility that this might be explicable as poor, rather than selective, imitation. However in this case, such an explanation seems most unlikely. We know through experiments with the pin-apple of the kind described earlier that in certain contexts both children and chimpanzees can and do imitate the particular sequential order in which they see a series of actions performed. Accordingly, we seem to have true selective imitation in these studies, with both chimpanzees and children.

Why might this be? We suspect, in part by intuition in observing the modeled actions ourselves, that the subjects may find it relatively transparent that the left-right order of operation is not critical; one can "just see" that it's not going to matter whether one starts on the left or right. By contrast, it may be a more sophisticated operation to recognize that there is an alternative solution to the particular, efficient-looking hierarchical alternative each subject sees. However, this is speculation; and we must remember that in actuality, both the sequential and hierarchical alternatives are arbitrary, so that in purely logical terms there is no reason to copy one and not the other. The device *might* have been built so that its internal structure meant that the particular chain of actions shown *was* the only one that would work. Hence these experiments, in concert with those described previously, leads us directly to a study specifically designed to test whether our speculative interpretations of selectivity are on the right track.

11.3 Selectivity of Imitation Contingent on Recognizing Causal Structure

The interpretations of earlier findings reviewed earlier are not the only foundations for the experiment outlined here. Another important one is a study by Want and Harris (2001 and vol. 2, ch. 8 by Harris and Want). Want and Harris presented preschool children with a horizontal tube containing a reward that could be pushed out using a stick. However, if the reward were pushed in one direction, it fell into a trap; only pushing from the alternative end constituted a correct solution for this task. In one condition of the study, the children were presented with a demonstration of the correct approach; in another, they were first shown an incorrect approach (pushing the reward to the edge of the trap, accompanied with the exclamation "Whoops"), followed by a correct one. Want and Harris found that 3-year-old children were subsequently more proficient at the task themselves if they had seen the incorrect-then-correct model, than the more straightforwardly correct one. Moreover, these children did not copy the incorrect part of the demonstration they had seen. Accordingly, we have here an apparent case of selective imitation. In the terms of our figure 11.1, the children selected out the causally irrelevant part of the incorrect-then-correct demonstration, presumably because their understanding of causality was sufficient to allow them to recognize it as such.

One of us (VH) conducted a study using 3-year-old children and young chimpanzees that incorporated some principles similar to those of Want and Harris, but had some extra twists. In this experiment, the subjects watched a model "stab-fishing" in one of two identically structured boxes. One of these boxes was opaque. The subjects watched as a model took a stick tool, used it to remove a bolt covering a hole in the top of this box, then stabbed the stick down through the hole three times (figure 11.5a). The stick was then removed; a small door on one side of the box was opened; and the stick was finally poked along the tunnel so revealed to fish out a small item of food (figure 11.5b).

The reader might like to imagine how he or she would perform the task given the opportunity. Our prediction is that the approach would be different than that elicited by observing the other version of the task, in which the box is transparent, except for the tunnel concealing the food reward. In this condition, an observer can see that when the stick is stabbed down through the top hole, it merely strikes a false ceiling and does no useful work in relation to the tunnel containing the food (figure 11.5c). We reasoned that after having seen this transparent version of the task, a chimpanzee or child might be more likely to omit the stabbing action in the top hole, assuming they could recognize it as causally irrelevant. In other words, they would veer to the emulation pole of copying shown in figure 11.1. By contrast, in the opaque condition there is less basis for recognizing the stabbing action as irrelevant, and we predicted that here a

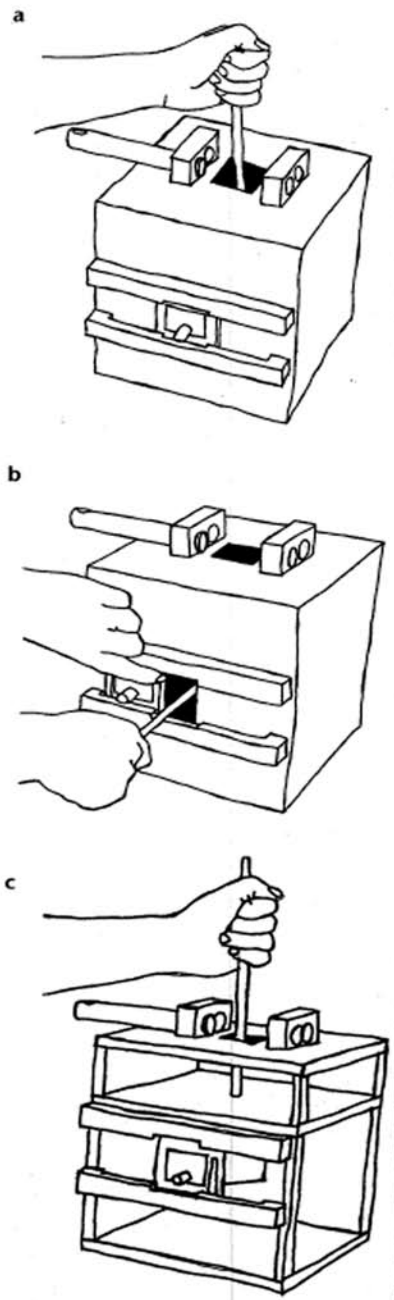

more fully imitative copy would be adaptive, and more likely to be performed by both child and chimpanzee.

The key results are summarized in figure 11.6 (Horner & Whiten, 2004). Three-year-old children clearly did not conform to our prediction. Instead, they largely imitated the whole routine they had witnessed, including stabbing in the top hole first, even when they had seen the transparent version of the model. This was true whether they saw the opaque version first or the transparent one first. It even occurred when the children were left alone to attempt the task. The children were not encouraged to imitate in any way, but merely invited to attempt to reach the reward themselves.

By contrast, the chimpanzees did conform to our prediction. Although with the opaque version of the task they did not copy as faithfully as the children did, they clearly implemented a different, more emulative strategy after observing the transparent version of the task modeled. Whether they saw this first or after a series of trials with the opaque version, they were significantly more likely to leave out the irrelevant action in the transparent condition than in the opaque one. We believe this is the first empirical demonstration of such an effect in nonhuman primates and now discuss its implications.

11.4 General Discussion

11.4.1 Adaptive Use of Imitation versus Emulation by Chimpanzees

Most discussion of the evidence for emulation and imitation has focused on whether a particular study of a particular species shows it to be emulating *or* imitating. Implicit in many such analyses is that that species might be characterized as showing one of these, but not the other. For example, Byrne and Tomasello (1995) argued that the claim by Heyes et al. (1992) to have shown imitation in rats might instead result from the rats being emulators. Likewise, Tomasello (1996), reviewing several experimental studies with apes undertaken up to that date, concluded that apes that had

◄ **Figure 11.5**

Task for studing selective imitation. (*a*) Once the top bolt is moved out of the way, the stick is stabbed repeatedly into the top hole of the opaque box. (*b*) The bottom door is then opened (here by sliding to the left) and the stick is used to fish out a reward from a tunnel. (*c*) When the transparent box is used, an observer can see that the stabbing action in the top hole is ineffective, for it terminates on a transparent false ceiling above the reward location.

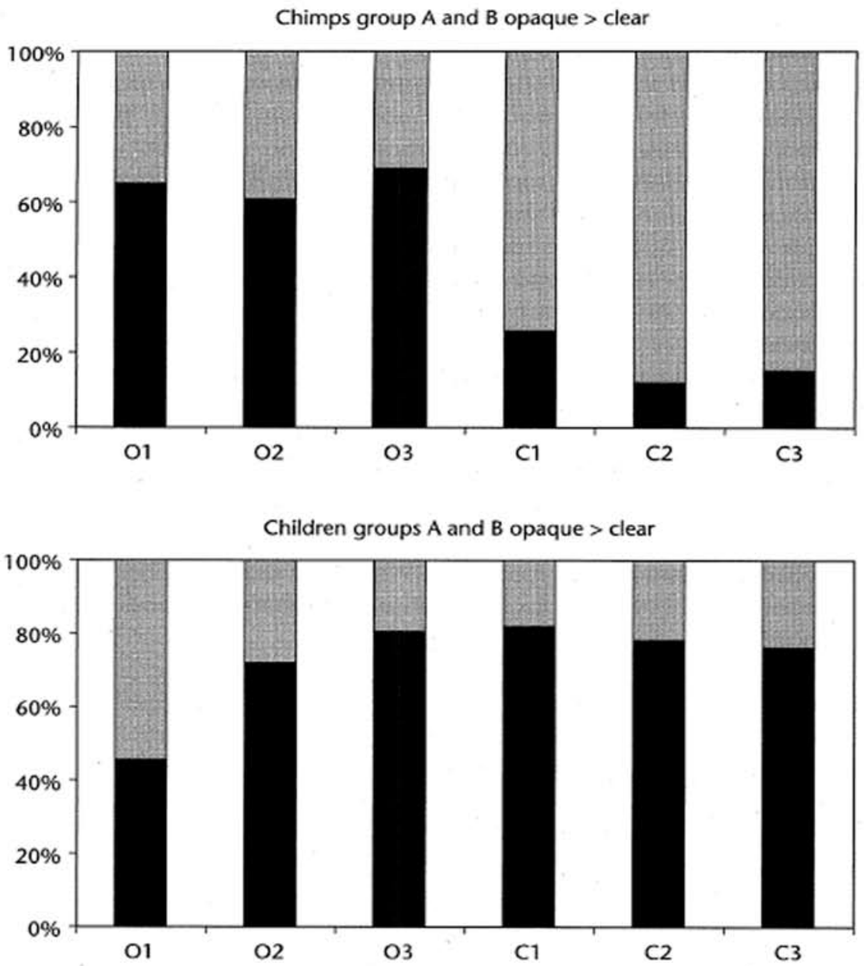

Figure 11.6
Results of the tests for selective imitation, showing percentage of actions first directed to the irrelevant top hole (dark portion of bar) versus causally relevant bottom hole (light portion). (*Above*) Chimpanzees showed a preference for starting with the irrelevant action on the top hole (dark portion of bar) in three trials of the opaque condition (O), shifting, as predicted, to focus more on the final, causally relevant action on the bottom hole once they were switched to the clear (C) condition (other subjects who started with the transparent box showed similar behavior). (*Below*) Young children did not show the shift in selective imitation seen in the chimpanzee subjects; in all conditions, they tended to copy both causally relevant *and* irrelevant components of the task.

not been exposed to human enculturation could be best characterized as emulators and not imitators. Our stab-fishing experiment suggests instead that for chimpanzees at least, aspects of the information available in the model may be selected in (a more fully imitative response) or selected out (a more emulative response) according to the subject's appraisal of the significance of this information. In our experiment, "significance" concerned the perceivable, noncausal role of the stabbing action in the top of the transparent box. Further studies will be needed to discover if similar flexibility in social learning applies to other information that might be used to guide the selectivity of copying, such as whether certain components are accidental or intended, an issue already addressed in the case of children (Meltzoff, 1995; Carpenter et al., 1998a; 2003) and beginning to be addressed in primate studies (Myowa-Yamakoshi & Matsuzawa, 2000).

It will be evident that these results lead us to believe the answer to the question, Do apes ape? (Tomasello, 1996) is "yes," but only insofar as this will depend on context. "Aping" seems traditionally to carry a connotation of relatively unintelligent mimicking, with little understanding of the meaning of the act copied. Our data suggest that instead, apes may copy intelligently; they will appraise the "meaning" of components of an act they see associated with desirable outcomes, such as those perceived as causally connected with such outcomes.

We do not wish to overinterpret what the chimpanzees are doing in making the discrimination about "causality" that our results confirm. Povinelli (2000) has reported extensive suites of experiments designed to find out just how much chimpanzees understand about causality. From these he concluded that chimpanzees take into account causality that can be directly perceived, but they do not appreciate underlying physical principles that are unobservable, such as force, which we humans do. Our results are not inconsistent with such conclusions, for the sense of causality that swung the imitation-emulation pendulum for our subjects hinged only on recognizing that (1) the transparent false ceiling blocked effective connection with the food location and/or (2) the stick had to make contact with the food location to have an effective role, two factors we have checked in separate experiments.

In this connection, it is worth putting our results in context by referring to some important bird studies, where we find some of the few other experiments to have examined selectivity in social learning. Palameta and Lefebvre (1985) showed that pigeons would learn by observation to peck through paper covers over food dishes, but they would do this only if the models they witnessed had obtained food in this way, in contrast to models that were observed pecking but gaining no food. Akins and Zentall (1998) have shown a similar effect in Japanese quail. There is a parallel with our own study in that these birds are showing adaptive selectivity, that picks out others' actions that are effective in gaining a desirable outcome. However, although the birds also detect causally relevant actions in this way, they may be doing so only at an associative level ("perform acts seen in others that are associated with feeding"). Our own results cannot be interpreted in this way. Actions on both opaque and transparent boxes were correlated with food reward, as was the stabbing action in both cases. The discrimination made by the chimpanzees must have relied on the cognitive appraisal of causality, at least in terms of the relevant connectivity of tool and food location.

11.4.2 "Blanket" Copying by Children

In our stab-fishing experiment, the children differed from the chimpanzees in copying in much more "blanket fashion," persisting in copying the irrelevant actions even with the transparent box, and even when the model left the room to remove social pressure for conformity.[2]

Two kinds of interpretation of such results can be entertained. The first acknowledges that this is not a freak result, but in fact mirrors many others provided by our own and others' research. For example with the pinapple task, as noted earlier, the children tended to copy much more of the task than the chimpanzees, even though they might, as a result, have taken longer to complete the whole task. We are thus led to recognize a real tendency of children, at least in contrast to nonhuman primates, to "overcopy." A plausible explanation for this is simply that we are such a thorough-going cultural species that it pays children, as a kind of default strategy, to copy willy-nilly much of the behavioral repertoire they see enacted before them. Children have the longest childhoods of any primate, much of which is spent in play, practice, and exploration, so there is plenty of opportunity to weed out wrongly assimilated aspects of the actions observed. Overcopying coupled with play might even provide a measure of serendipitous innovation in the culture at large.

A second interpretation acknowledges that there are other findings in the literature that do show selectivity; for example, those of Want and Harris (2001) and Carpenter et al. (1998a) have already been alluded to. So of course we would not wish to claim that our result show that children never exhibit selectivity. Even with our specific task, it becomes an important question as to whether older children—or perhaps even younger ones—will see the social world differently and copy in a less blanket fashion. One factor that may well explain why Want and Harris and Carpenter et al. obtained different results is that the models in their studies marked the irrelevant actions with a vocalization, like "Whoops." Another factor might be that we repeated the demonstrations several times, even before a child made its first attempt. This may have been construed by the child as indicating that the stabbing action was indeed important, perhaps because it appeared necessary rather than inconsequential, or because it was intended rather than accidental. This line of thinking leads to the intriguing possibility that the children copied all they saw done on the transparent box, not because they were less inclined to be selective than were the chimpanzees, but because they were also actively assessing the meaning of what they saw, yet unlike the chimpanzees saw stabbing as a repeated, therefore intended, and therefore likely important component of the task. Further experimental manipulations would be required to examine such possibilities.

2. We invite the reader to imagine what kind of interpretations would have been offered by themselves and various theorists in the social learning community if the child and chimpanzee results had come out the other way round!

11.4.3 Imitation of Sequential and Hierarchical Structure

In discussing our experiments designed specifically to assess the imitation of these aspects of the structure of actions, we noted selective assimilation of hierarchical structure while aspects of sequencing were ignored by both children and chimpanzees. However, the stab-fishing task also casts light on this issue and in fact shows the opposite effect: copying of both hierarchy *and* sequence. Our reasoning about this rests on the conjunction of two of the findings: (1) that in one condition the subjects performed the whole task, but in the other only the second part of it; and (2) that in the full performance they carried out the sequence of actions they had witnessed (i.e., stab in the top hole, then fish in the side hole). Finding (1) shows that the chimpanzees could decompose the task (parse it, in Byrne's, 1999, linguistic analogy; see also vol. 1, ch. 9) and thus perceived its hierarchical structure, and finding (2) shows that they recognized and copied the correct sequential structure when they performed the whole task. One of us has argued elsewhere (Whiten, 2002a) that one reason the concept of program-level imitation (Byrne, 1994) may be misleading is that it conflates what are in principle (and, it seems, in practice) the two separable issues of imitating sequential versus hierarchical structure. However, to the extent that the definition of program-level imitation includes copying both sequence and subroutine structure (R. Byrne & Russon, 1998, p. 677), we believe the stab-fishing results provide the first quantitative evidence for this particular conjunction in a nonhuman species.

Why do we find copying of both sequential and hierarchical structure in stab-fishing, but not in the other tasks we used to investigate this issue? We suggest that the most likely interpretation hinges on the arbitrariness built into the two-action experimental design. In the case of the key-way fruit, for example, it does not matter whether the task is started on the left or the right, and both children and chimpanzees appeared to recognize this. In the stab-fishing task, however, if one recognizes the stabbing as an element to be copied, it can only be done sensibly before the fishing element or it is redundant. Of course, the nature of the experimental design means that in reality, whether the stabbing or the fishing is done first *is* arbitrary to the extent that a subject may perform these in either order, i.e., stab then fish, or fish then stab (and note that likewise, in the key-way task, in reality there are numerous ways to complete the task, yet the subjects tended to use the one demonstrated). Thus the word "sensibly" is important in the earlier sentence; we are not appealing here to arbitrariness in the reality of the task as set by the experimenter, but rather we are making a suggestion about how the subjects are construing the task, and more particularly, why they might rationally apply different rules of selective imitation in each.

11.4.4 Extending the Two-Action Methodology

We end with a methodological point, which given continuing controversies in this field (Caldwell & Whiten, 2002), seems worth emphasizing. The principle of the two-action task has been advocated by many in the field of social learning to be an important tool in discriminating forms of social learning, imitation in particular (e.g., Galef et al., 1986; Heyes, 1996; Zentall, 2001). By contrast, Byrne (2002a) has argued that the approach fails to identify imitation because the tasks are insufficiently novel to do so. It is true that the two-action method has mainly been celebrated by those working on birds and nonprimate mammals and using relatively simple action alternatives, such as pushing a bar to the left or the right, or pecking versus stepping on a treadle. However, this is not an intrinsic limitation in the concept of the two-action method. If novelty is the problem, then it just requires that the two alternative actions modeled (or minimally, perhaps at least one of them) be shown to be "novel" according to some agreed-upon operational criterion, such as that provided by a baseline, no-model condition. More important, the two-action approach need not be limited to a too-myopic definition of action, which limits the alternatives

to relatively small differences in the type of bodily actions (e.g., push left versus push right), as Byrne appears to assume.

To the contrary, we hope we have shown that the fundamental approach can be applied with rigor and with interesting results in relation to higher levels of complexity, such as those of sequential and hierarchical structures, and the availability of causal information. In principle, the approach appears to be a powerful way to investigate selective imitation in relation to just about any contrasting aspects of actions, however complex, that can be copied and that we psychologists can imagine.[3]

Acknowledgments

We thank the Biotechnology and Biological Sciences Research Council for funding the recent research by SM-P and VH described here. We are grateful for the exceptional support of Debby Cox, Cherie Montgomery, and the trustees and staff of Ngamba Island Chimpanzee Sanctuary and Wildlife Conservation Trust, where these studies were conducted.

3. See comments on this chapter by Galef vol. 1, ch. 12.5, p. 295, and by Jones ch. 12.6, p. 297. For relevant discussion see chs. 6, 16, 17, and 18 in vol. 2, by Harris and Want, Gil-White, Greenberg, and Chater, respectively.