Part III

'Actualizing' Conditions for Innovation in Stone Knapping

Chapter 20

Transfer According to the Means in Human Infants: the Secret to Generative Tool-use?

Emily W. Bushnell, Jason Sidman & Amy E. Brugger

This paper focuses on the cognitive skills and abilities involved in tool using and tool making, and an analogy is drawn between the behaviour of developing infants and scenarios which may have occurred among hominins 2–3 million years ago. It begins with a review of the developmental and comparative literature on several cognitive capacities central to making and using tools, including the ability to project into the future, the ability to differentiate and relate means to ends, and the ability to imitate object-directed skills. Original research with human infants is presented which builds on this prior literature and establishes that infants attend to the model's causal actions in demonstrations of means–ends behaviours and furthermore are able to transfer these actions learned by imitation to novel circumstances. An important implication of these results is that such transfer increases the probability of discovering useful new means–ends behaviours even without full causal understanding or 'insight'. Thus transfer of the means may be a significant source for invention of new behaviours during an intermediate phase of causal cognition. The paper closes with the speculation that a parallel process of transfer according to the means could also have operated in evolutionary time and been a part of how human technology progressed from the original case of stone knapping.

Making and using tools requires a unique confluence of anatomical features, neuro-motor skills, intellectual capacities, and social and environmental circumstances. Many of these prerequisites are addressed in the various contributions to this volume, which targets the prototypical instance of tool use and manufacture, knapping stone. As psychologists, our focus and expertise is on the cognitive skills and abilities demanded by such activities. For example, to make and use even the simplest kind of tool requires at least intention, memory, learning, foresight, attention, and sequencing (Schick & Toth 1993; Wynn & McGrew 1989). More advanced technologies may additionally require mental imagery, categorization, decision-making, evaluation, and spatial concepts such as symmetry and balance (Corballis 1999). Any early hominin or proto-human who systematically made and used stone tools must therefore have possessed these mental faculties, and thus archaeological finds of stone tools can serve as benchmarks helping to characterize the evolution of intelligence.

The cognitive prerequisites for using tools have also been emphasized in analyses of child development. The logic, as for the hominin above, is that in order to effectively exploit objects in particular ways, a child must possess the cognitive abilities entailed. Therefore, the development of tool using in individuals will be paced by cognitive development, among other things, and conversely what children are able to do with objects can serve to diagnose their level of cognitive maturity. This reasoning is the basis for several standard assessments of infant growth and development, for example the Uzgiris-Hunt Ordinal Scales of Psychological Development (Uzgiris & Hunt 1975) and the Bayley Scales of Infant Development (Bayley 1969).

The difference between the early hominin and the child, of course, is that the child is here and now and can be confronted empirically with various object tasks and situations. Suites and sequences of abilities with objects can be firmly established and conditions which promote these can be identified. These developmental findings in turn may be useful to scientists working to define how technologies such as flint knapping first arose in humans and why they exist only in humans. While we do not think that ontogeny necessarily recapitulates phylogeny, the cognitive problems posed by acquiring and executing object-directed skills and the potential set of solutions to these challenges are similar in many respects for children and early hominins. Like infants in particular, early hominins presumably did not have access to symbolic instructional media (i.e. language, diagrams, etc.), but they did operate in a rich social milieu and in an environment 'cluttered' with objects. Furthermore, given that the 'final state' for both the developing child and the evolving hominin is the same, namely the cognitively-mature modern human, there are only so many intermediate states of cognition that are logically possible. Thus, those which can be empirically documented in the case of the developing child are worth considering as perhaps applicable to the evolving hominin as well.

Drawing out this analogy between the cognitive status of developing infants and scenarios which may have occurred among hominins 2–3 million years ago is our general strategy in this chapter. We will first briefly review the infant and comparative literature on several cognitive capacities central to making and using tools, including the ability to project into the future, the ability to differentiate and relate means to ends, and the ability to imitate object-directed skills. We will then present some of our own research with infants which builds on this prior literature and examines the phenomenon of generalization or 'transfer' of imitated actions. Finally, we will elaborate on an important implication of the results of our research, and outline how a parallel process might conceivably have operated in evolutionary time. Our goals are to advance what is known about infants' object-directed skills and the development of tool use, and with that knowledge to provoke healthy speculation about how and why human technology may have realized and then progressed beyond the original case of stone knapping.

Cognitive capacities central to making and using tools

Future-oriented thinking

We begin by noting that both tool using and tool making are clearly means–ends behaviours. In each case, an action or sequence of actions is performed at one time in order to bring about a particular product or event moments, hours, or even days later. An important cognitive component of means–ends behaviour is future-oriented thinking or 'foresight', the ability to run time forward mentally and 'see' an end state before it has occurred. Indeed, looking ahead in the mind to a particular end state, that is needing it or wanting it, may be thought of as the stimulus which prompts performing the initial action or means in the first place. In some sense, any operant behaviour emitted to obtain a reinforcement involves future-oriented thinking, but with tool using and tool making, the ends do not necessarily follow the means so immediately or directly; large spans of time and numerous other events may intervene. With early stone knapping, for example, a certain type of strike (at a given angle, with a given force, etc.) may have been prompted by the desire to have a certain edge just moments later, but selecting the core with an eye to producing that same edge may have occurred at a remote site and thus hours or more in advance. Similarly, producing the edge and perhaps carrying it around would have been predicated on the intention or need to use it at some point in the future well beyond the next instant, perhaps even at some indeterminate point in the future.

This parameter of the span of time to which one can look ahead has been emphasized by both anthropologists and developmental psychologists. For example, Boesch & Boesch (1983) observed that chimpanzees in the Taï forest carry nuts and stones with which to crack them open to anvil sites as far as 500 metres away and often out of sight from the starting place. Because transport distance equates to travel time, this behaviour implies a certain ability to look ahead to a future ends or goal-state. Meanwhile, Visalberghi (1993) found that capuchin monkeys did not traverse even to an adjacent room to fetch a tool for obtaining food at the starting place, implying a rather short time span for planning. In comparison, Wynn & McGrew (1989) noted that the makers/users of Oldowan tools carried stone to work-sites as far as 6 kilometres away from its source, implying perhaps a longer time span of prospective thinking than the chimpanzees show. Of course there are other factors involved in these behaviours as well, such as spatial-memory capacities, the natural ecology in which different species live, the availability of alternative food sources, and so forth, but if we set these aside or presume them constant for a moment, a pattern seems to emerge. Namely, the closer a primate species is to modern-day humans, evolutionarily speaking, the longer the span of future-oriented thinking it shows. The makers/users of Odowan tools could apparently plan further ahead than chimpanzees, who in turn can evidently plan further ahead than capuchins.

The same pattern of an expanding time frame for planning is observed over the course of early development. For example, at the simplest level of thinking ahead over a matter of a few seconds, Canfield & Haith (1991) found that 3-month-old infants learned to anticipate the reappearance of visual stimuli over a greater number of steps and thus a longer time frame than 2-month-olds could anticipate. Similarly, with a more complicated task involving manual skills, Willatts (1999) found that 7-month-olds were able to intentionally pull on a cloth in order to retrieve a toy resting on it, provided the toy was at the near end of the cloth and thus could be retrieved immediately with the initial pulling act. However, 8-month-olds intentionally pulled on the cloth even when the toy was at the far end and thus required several tugs and more time to reel in. With a related pull-to-obtain task, Boudreau & Bushnell (2000) likewise documented that older infants can 'look ahead' through several steps and over a time-frame of about 10 seconds as they commenced a manual task.

Perhaps the clearest demonstration of the developmental increase in time span which can be accommodated between means and ends comes from some clever work by McCarty *et al.* (1999). These researchers examined how infants approached a spoon loaded with food which was placed at different orientations in front of them. They observed that with older and older infants, adjustments in gripping the spoon to successfully get food in the mouth were made progressively earlier and earlier in the sequence of actions. The youngest infants, 9-month-olds, frequently grasped the bowl of the spoon and only switched to grasp the handle after they had put the handle in the mouth and received no food. Fourteen-month-olds adjusted their grips 'mid-stream'; they grasped the spoon's handle but sometimes with the thumb pointed away from the bowl and then shifted the spoon to the other hand to achieve the optimal grip as they moved the spoon to the mouth. Finally, 19-month-olds selected which hand to use or whether to grasp the handle over-handed or under-handed right at the start of the trial, so that they could put the spoon in the mouth smoothly from there, with no further adjustments. In short, McCarty *et al.*'s results showed that with their initial approach to grasping the spoon, older infants planned further ahead toward the aim of feeding themselves than younger infants.

Thus in the context of several different tasks — anticipatory looking, pulling one thing in to obtain another, and using a spoon to transport food — developmental psychologists have found that older infants were able to look ahead across more action steps, greater distances, and longer time intervals than younger infants. This developmental pattern parallels the one observed earlier within primate and hominin evolution. In both cases, as the human species evolved and also as the human infant develops, the cognitive ability to run time forward from an active means to a desired ends seems to expand, and individuals can look ahead farther and farther into the future to regulate concurrent manual behaviour. Similar changes in brain anatomy and functioning might be responsible for these parallel improvements in future-oriented thinking, or it is also possible that altogether different factors mediate the evolutionary and developmental trends. Either way, though, the types of tool-using and tool-making skills logically possible at particular values of this looking-ahead parameter would be similarly constrained. Thus, as we argued earlier, it may be instructive for understanding the underpinnings of seminal behaviours such as stone knapping to examine the tool-using and object-manipulation abilities of infants.

Understanding physical causality

Comparative and evolutionary considerations

Even more than the span of time over which ends can be seen from their means, anthropologists and developmentalists have emphasized appreciating the *connection* between means and their ends as a critical cognitive component for object-directed skills and tool using. As Tomasello (1999) clarifies, anticipating what end state will follow a particular means is not the same as knowing 'why' that particular means leads to that end state, that is, appreciating the causal forces and spatial relationships which ensure that it will be so. Although the former can be learned from repeated experiences, the latter may facilitate faster learning, the repair of failed attempts, the construction of new or substitute means in the absence of the initial one, and so on.

Accordingly, anthropologists have focused on causal understanding as an important distinction between species which use tools flexibly and frequently and those which use tools rigidly, in narrow contexts, and in highly-rehearsed ways. For example, although they readily learn to do so in captivity (see Cummins-Sebree & Fragaszy this volume), capuchin monkeys rarely use tools in the wild, whereas chimpanzees use a variety of tools and some 15 per cent of their food processing in the wild involves tool use (Matsuzawa 1994). Laboratory research also supports this cognitive distinction between apes and monkeys. In their extensive review of the experimental literature on primates and tool using, Tomasello & Call (1997, 95) concluded that 'apes may have a deeper causal understanding of tool-use tasks than monkeys'. Byrne (1995, 98) likewise concluded that 'all the great apes, but probably no other animals, can comprehend the simple cause-and-effect relationships that govern the use of objects to solve problems'. Most telling in Byrne's analysis is the variety of errors that capuchins made when confronted with novel variations on a task they had mastered. They used sticks which were too short and too thick to poke a peanut out of a transparent tube, for example, whereas chimpanzees solved these extraction tasks insightfully (without trial-and-error). Importantly, the solutions to these tasks required the actor to *create* a certain spatial relationship; with pull-to-obtain problems where the presence or absence of the necessary spatial support is visible, capuchins as well as chimpanzees performed successfully. Thus the ability to mentally represent relations between objects may be a critical difference between monkeys and apes so far as causal understanding is concerned.

When multiple relations must be represented, however, even apes seem to be challenged. Byrne (1995) points out that despite the range of tools they use in the wild, chimps do not make any tools by adding one item to another (as a shaped stone and a stick are added to make a spear, for instance). Likewise, although they do modify objects to use as tools, they do not make any tools which are used for making another tool (as humans make knifes in order to sharpen digging sticks, for example). Similarly, Tomasello & Call (1997, 95) note that when a trap component is added to the tube-extraction task mentioned earlier, thus requiring an additional relation to be represented, 'apes are far from perfect, (and) require many trials for learning in some variations'. Furthermore, once they learn the appropriate strategy for avoiding the trap, they persist in using it even when the spatial arrangement is changed and no longer demands it. Mature modern humans solve all of these tasks, of course (they invented them!), and so the ability to represent not just object relations but sequences and hierarchies of object relations may be a critical difference between apes and humans with regard to causal understanding.

Thus, as with the time-span for prospective thinking, we again see a sort of scaling from monkey to chimpanzee to human, so far as causal understanding goes. Monkeys (in particular capuchins) do not seem to appreciate mechanical, cause-and-effect relations, except perhaps when they are transparent in the visible layout of objects. Apes (in particular chimpanzees) may appreciate simple causal relations, but are not able to manage sequential or embedded strings of them, which humans in turn can master. Where on this rough continuum the earliest stone knappers might lie is something of an open question. Obviously, cloth-pulling and tube-extraction tasks cannot be posed to extinct species, and there is disagreement in the literature concerning the level of understanding with which the earliest Oldowan artefacts were made. Some have pointed to the lack of symmetry and retouching in these tools as an indication that there was no real intended design; thus Wynn & McGrew (1989) and Toth (1985) argue that Oldowan tools require no more than the ape 'grade' of intelligence. On the basis of more recent evidence, however, others argue that even Oldowan tools betray multi-step planning and complex intentionality in their making (see Roche this volume). Still others have questioned whether the heralded Acheulean handaxe really represents any cognitive advance over the earlier Oldowan tools (see Davidson & Noble 1993). Without the opportunity to observe these ancient artefacts in the making, these issues may be very difficult to resolve. However, it seems fair to say that at some point in evolutionary time between now and the days of our common ancestor with apes, there was a level of causal understanding intermediate to those of modern chimpanzees and modern humans.

Understanding physical causality: developmental considerations

As a window on what such an intermediate level of causal understanding might be, we now turn to the field of child development, where an incremental achievement of causal understanding has also been examined and outlined. Indeed, many of the tasks used in the primate studies cited above were employed first to test babies and children. The seminal work here is that of the pre-eminent developmental psychologist, Jean Piaget (1952; 1954), who thoroughly dissected the means–ends abilities of infants and charted their emergence and development through careful observations of his own three children during the first two years of life. According to Piaget, an important transition occurs at about 8 months of age, when infants begin to differentiate the means from the ends in their causal activities with objects. Prior to this transition, these activities are experienced as unanalyzed and unordered wholes. Seeing the effect (e.g. instigated by someone else) is as likely to provoke the infant to act as the infant's action is to create the effect, in a phenomenon Piaget labelled the secondary circular reaction. In novel situations, infants' haphazard behaviour may fortuitously produce effects, and they may then experience a sense of efficacy and systematically repeat the precipitating actions. However, Piaget argued that even while they deliberately repeat effective actions, young infants are unaware of the particular spatial and temporal relations which mediate between their actions and an effect. Thus their sense of causal agency is largely 'phenomenalistic' or magical, and their efforts to recreate effects may often prove to be irrational and ineffective. A true understanding of causality, according to Piaget, is implied only when the infant intentionally acts in a deliberate fashion so as to reliably and efficiently obtain a desired outcome. This level of understanding is betrayed by premeditated actions before an effect has been observed, for example, and by systematic searching for a cause upon observing an unprompted effect.

The developmental difference described by Piaget (1952; 1954) and his emphasis on intentionality or prior expectation are supported by the results of Willatts's (1999) empirical work with cloth-pulling, mentioned earlier in the context of prospective thinking. In addition to the distance effects discussed there, Willatts also found that although 6-month-olds and 8-month-olds were equally likely to retrieve a toy by pulling in the cloth it was on, their behaviour was qualitatively different. Six-month-olds' retrieval of the toy was often an unexpected 'side-effect' of grasping the cloth to play with it (the cloth). In contrast to 8-month-olds, they did not maintain fixation on the toy as they pulled the cloth, and they showed longer time intervals between contacting the cloth and contacting the toy, filling this time with shaking the cloth, mouthing it, etc. before 'noticing' the now within-reach toy and grasping for it. Even more telling, 6-month-olds were just as likely to pull on a cloth when there was no toy on it as when there was a toy, whereas 8-month-olds pulled the cloth more frequently when a toy was there. For the older infants, pulling on the cloth was a deliberate means to an end (obtaining the toy), but for the younger infants pulling on the cloth was a behaviour in and of itself, which then fortuitously brought a toy within reach if one happened to be there.

Thus understanding the means to an end as such in effective activities is a developmental achievement that unfolds during infancy in humans. Moreover, the emergence of causal understanding is not an all-or-none phenomenon, but is affected by the same spatial transparency and embedding factors which are relevant in analyses of nonhuman primates' causal understanding. Both Bates *et al.* (1980) and also van Leeuwen *et al.* (1994) found that infants were more likely to systematically use various tools to obtain a goal toy if the tool was presented in spatial contact with the toy; only older children systematically used the tools when they were not initially in contact with the toy and the children had to generate this contact themselves. Similarly, both van Leeuwen *et al*. (1994) and Bushnell & Boudreau (1998) have discussed the parameter of 'complexity' in children's tool use and the understanding which mediates it. Bushnell & Boudreau, for example, discuss several instances where children can successfully control each component of a tool-using task, but are unable to perform effectively when these must be nested together. For the purpose of comparing both across-species and across-age (within humans) then, causal understanding appears to be very domain-specific; an individual can 'have' it with regard to one kind or level of problem and not with regard to others.

For any particular means–ends task, though, an important developmental question (and evolutionary one too, we suppose) is not so much when the transition to causal understanding occurs for an individual or a species, but how it occurs, that is, what drives the change that takes place. Piaget argued that, somewhat ironically, it is the very haphazard, unintentional behaviours of the pre-causal infant which blossom into goal-directed actions that qualify as genuine means to ends on the part of older infants. Through acting frequently and variably with objects, and attending to the consequences of these activities, infants become sensitive to the relevant spatial and temporal contingencies and internalize them as causal knowledge. Thus acting comes first; within any given task domain, infants initially behave so as to cause effects without knowing how or why their actions do so.

In contrast to Piaget, others maintain that infants understand aspects of causality well before they act either haphazardly or systematically in means–ends tasks. For instance, some researchers argue that very young infants understand principles of spatial support (Baillargeon 1994) and human agency (Leslie 1984; Woodward 1998); although they do not act on this knowledge, they evidence it by looking longer at visual displays which violate the principles in question. These researchers suggest that casual understanding is either present innately or derived very early in life through applying general learning mechanisms to numerous observations of everyday events and other people's actions on objects. Young infants perform unsystematically in tasks such as Willatt's cloth pulling because of limitations on attention, inabilities to inhibit pre-potent behaviours, and difficulties coordinating their knowledge with motor behaviour. As soon as these ancillary problems are resolved, infants' behaviour in means-ends tasks shows itself to be intentional and predicated on causal understanding which was present all along.

Although the debate just outlined continues, two recent direct comparisons of infants' performance on means–ends tasks and their independently-assessed understanding of the same tasks seem to support Piaget's original suggestion that successful acting precedes causal understanding. Schlesinger & Langer

(1999) presented some infants with several different object-retrieval tasks similar to those employed by Willatts (1999). Other groups of infants were shown visual displays of 'possible' and 'impossible' toy retrievals in the same task contexts. The researchers found that infants were able to perform each task effectively at a younger age than that at which they showed comprehension of the corresponding visual displays. For instance, as in Willatt's studies, 8-month-olds systematically pulled on a cloth to retrieve a toy when the toy was on the cloth but not when it was beside it. However, they did not also look longer at displays where a toy came forward even though it was not in spatial contact with the cloth; only 12-month-olds showed a reaction to these 'impossible' displays. The authors concluded from results such as this that within each task domain, causal action precedes causal 'perception' or understanding.

Woodward & Sommerville (2002) conducted a similar study but employed a within-subject design. They tested the same individual infants in both a cloth-pulling task and a looking paradigm designed to assess their interpretation of someone else's cloth-pulling. The results showed that only those infants who intentionally pulled cloths to retrieve toys themselves seemed to recognize the toy as the actor's goal in the looking paradigm. Conversely, infants whose performance in the cloth-pulling task was disorganized and not systematically related to the toy's position seemed to interpret the cloth as the actor's goal in the looking paradigm. These results along with those of Schlesinger & Langer (1999) suggest that infants do not understanding why pulling a cloth makes the toy obtainable, that is, they do not appreciate the spatial relation between the cloth and the toy, until after they are able to effectively pull cloths to obtain toys! Thus in development, causal understanding may be formulated in a *post hoc* fashion, derived from a kind of reflection upon or 'noticing' the outcomes of one's successful behaviours.

Understanding physical causality: a conundrum

The developmental sequence just discussed highlights a subtle but important distinction between 'intentionally' in the lay sense and 'with intentionality' in the stricter Piagetian sense. To whit, one can perform an action (the means) intentionally, that is purposefully and with every expectation that a certain result (the ends) will ensue, but still without Piaget's intentionality or true understanding of the mechanical relations which link the means to the ends. This is the intermediate position of the infants observed by Schlesinger and Langer (1999) and Woodward & Sommerville (2002), where the infants could perform effectively on certain problems but could not yet perceive effective solutions.

It is possible that this same seemingly backwards sequencing of causal understanding and performing means–ends behaviours might also apply across species or evolutionarily. Both Byrne (1999) and Tomasello (1999), for instance, conclude in their reviews of the primate literature that although chimpanzees certainly use tools purposefully and effectively (intentionally), they may not understand (have intentionality) the dynamic, physical relations among the objects in their instances of tool use. This failure to understand would prohibit chimpanzees from readily inventing new tools and uses for them, and indeed technological advances are slow to accumulate in primate cultures and require multiple trials to learn in the lab. Along these lines, Byrne (1995) has qualified the quintessential instance of causal understanding on the part of non-human primates, and suggests that Kohler's chimpanzee Sultan may have played at joining sticks together before the famous occasion on which he did so to obtain out-of-reach bananas. That is, Sultan's intelligent behaviour may have been a case of noticing after the fact of performing rather than genuine 'insight' or performing *de novo* predicated on causal understanding. Similarly, some have characterized the earliest tool making of hominins as '*ad hoc*' or executed effectively but without full understanding of the spatial and mechanical factors involved (Wynn & McGrew 1989; Mithen 1996), although this interpretation is vigorously disputed by others (Pelegrin this volume) and would be difficult to establish (or not) with only artefacts as evidence.

In any event, the possibility that an infant or an ape or even an early hominin might systematically and purposefully engage in a means–ends behaviour without understanding why or how it works seems paradoxical. Why would they do so? This state of affairs — observed in the case of infants, suspected in the case of apes, and posed here for the sake of argument in the case of early hominins — demands that *something other than* causal understanding must sustain and support the initial instances of systematic means–ends behaviour. If understanding comes from noticing and reflecting after doing, then the doing must be predicated on other grounds. What could such other grounds be?

One mechanism which could maintain effective behaviour without cause–effect understanding is operant conditioning. However, this is only plausible if the effective action or 'means' (unknown as such) is likely to occur spontaneously and then also 'happens' to lead to a positive outcome relatively frequently. These conditions are readily met for certain support relations. For example, grasping a cloth just to mouth or attain it will inevitably bring anything resting on the cloth closer and perhaps within reach for the infant. In contrast, the odds of finding a nut sitting on a rock and happening to strike it with another rock are much lower than those for grasping a colourful cloth. Similarly, it would require quite a rare confluence of events for breaking one rock with another to coincidentally be followed by a rewarding consequence and thus learned by conditioning processes.

As an alternative to conditioning, Lockman (2000; this volume) has suggested that many behaviours which ultimately become functional may originate in infants' explorations of objects' perceptual properties or 'affordances'. For instance, towards the end of the first year infants frequently engage in banging objects against surfaces, presumably for the pleasure of making noise. In the course of this activity, they become sensitive to properties such as the weight of objects and the compliance of surfaces. They also master the motor elements of hammering, becoming adept at aiming, at controlling the angle between the object and the surface, at varying the force of the stroke, and so forth. These skills are then available to be easily incorporated into functional hammering later on. Similarly, lateral hand movements initially practised in the context of exploring textures (see Lederman & Klatzky 1987; Bushnell & Boudreau 1991) may become incorporated into scribbling with crayons and ultimately into writing. In these examples, particular actions which eventually are involved in tool-using and means-ends behaviours are initially sustained by perceptual feedback. Once actions are perfected and objects become familiar in this context, the cognitive 'leap' to imagining them in functional roles may be minimal. Indeed, with such practice and familiarity, straight-forward causal relations may literally be 'seen' in a quasi-perceptual process, as opposed to worked out thoughtfully. This process is further described by Cummins-Sebree & Fragaszy (this volume), who propose it as the likely foundation for capuchin's solutions to simple obtaining problems.

Imitation of goal-directed activities

A final and very viable mechanism which could support early instances of means–ends behaviour without causal understanding is imitation, or more broadly, observational learning. Intuitively, children and adults alike routinely exploit many objects without understanding how or why they physically work, and almost always, they have seen someone older or more knowledgeable demonstrate the activity beforehand. Operating complex machines such as cars, telephones, and microwave ovens come to mind as cases in point. Empirically, there are many reports in the literature of children learning means–ends behaviours by imitation. For example, with regard to tool using in particular, Bushnell & Boudreau (1998) and also Chen & Siegler (2000) found that although toddlers hardly ever spontaneously tried to use novel tools to obtain out-of-reach toys, they readily and successfully did so after seeing a model demonstrate using the tools. Similarly, van Leeuwen *et al.* (1994) found that following a demonstration, children were able to succeed with more complex tool/goal configurations than they could manage spontaneously. These results highlight the 'added value' of imitation for acquiring skills, above and beyond whatever exploratory tendencies and elementary causal knowledge children might bring to the task themselves. Even preverbal infants can imitate highly unconventional, novel actions toward objects, actions which they never exhibit while exploring the objects on their own (Meltzoff 1988). Indeed, imitation is so pervasive among children and at such early ages that Meltzoff & Moore (1989; 1994; 1999) have proposed an innate 'Active Intermodal Mapping' mechanism which permits even newborns to relate postures and movements they see in others to their own proprioceptively-experienced postures and movements.

The importance of imitative abilities to tool use has also been acknowledged in conceptions of the evolution of culture and technology. For example, Tomasello (1999) has emphasized that imitation permits new strategies discovered by especially intelligent individuals or through rare and fortuitous circumstances to accumulate, distribute, and be maintained within a group and across generations. He calls this important consequence of imitation the 'ratchet' effect; there is no slippage with the loss of knowledgeable individuals. Even more to the point, Merlin Donald (1991) has argued that a broad capacity to re-enact events was the first major evolutionary transition that set the cognition of hominins truly apart from that of apes some two million years ago. Among other things, this 'mimetic' ability allowed for the perfection and transmission of tool-using and tool-making skills, even in the absence of symbolic communication systems (language).

To Donald's position, some would immediately counter that other animals and especially non-human primates are capable of imitation. Certainly to the extent that some present-day apes use tools and other complex manual skills in their foraging, juveniles seem to acquire these abilities through observing adults execute them. Matsuzawa (1994) has described in some detail how young chimpanzees learn to crack nuts with stone hammers in the context of adults engaged in nut cracking, and Byrne (1999; this volume; Byrne & Byrne 1993) has similarly analyzed how young gorillas learn to manipulate particular plant foods from observing their mothers feeding. However, in the same breath these researchers also grant that the imitative abilities of nonhuman primates are qualitatively different from those of human infants. For instance, Byrne (1999) stresses that apes have usually seen repeated demonstrations — that is, in the hundreds — before they begin to copy a modelled action sequence. In contrast, humans including infants often successfully imitate object-directed actions after just a single demonstration. Byrne suggests that ape imitation may be mediated by a statistical learning process that involves motor priming; human imitation on the other hand may be mediated by more abstract mental representations and understandings of actions.

Want & Harris (2002) have recently provided another approach to distinguishing different qualities of imitation. In a comprehensive review aimed at both comparative and developmental psychologists, they scrutinized imitative behaviour and decomposed it into smaller constituent components. Following other analyses, they note that a given instance of reproduced behaviour could be the result of the imitator's attraction to the stimulus itself (stimulus enhancement), the imitator's attraction to the location of the observed behaviour (local enhancement), the reproduction of only the means (mimicry) or of only the ends (goal emulation) without consideration of the other, or the intentional linking of an observed means to achieve a desired end (true imitation). Want & Harris close with the suggestion that a set of criteria and also further research based on these separable components may facilitate both across-species and across-age comparisons with regard to imitative abilities.

Some researchers have already begun to evaluate imitative behaviours in light of Want & Harris's scheme. For example, Tomasello (1999) thinks a critical difference between apes and humans is that human children attend to and duplicate the model's precise actions or 'means' in means–ends behaviours, while apes only seem to learn what end results are possible from demonstrations. Thus, strictly speaking, apes do not imitate but instead only 'emulate' the model by recreating the observed effect in their own way. Tomasello and colleagues have documented this distinction in several studies directly comparing the behaviour of chimpanzees and human children when both were shown means–ends activities such as retrieving toys with a rake or a reel (Tomasello et *al.* 1993; Nagell *et al.* 1993). Children (and also enculturated chimpanzees) were more likely to imitate the experimenter's methods than mother-raised chimps were; the chimps often employed alternative means and either achieved the ends by some different route or failed to achieve the ends.

Tomasello (1999) and also Meltzoff (1995) have identified another way in which ape and human imitation may differ as well. Here they maintain that human children are uniquely successful at imitating because they are able to appreciate other people's intentions as they manipulate objects. Children inherently understand that other people are 'like me', and thus predicate their own manipulations of objects on the assumption that objects will do 'for me' what I have seen them do 'for' other people, if I act on them in the same way. Bard & Vauclair (1984) have also emphasized a social distinction between apes and humans with regard to object manipulations and the potential for imitation. In some direct comparisons, they found that in contrast to human adults, ape adults did not attend to their infants' activities with objects or attempt to facilitate and structure them in any way.

Any or all of these differences in the imitative process — concerning the number of demonstrations required, concerning whether the means as compared to the ends is attended to, and concerning the social aspects of the imitative context — may be important to technological differences between apes and humans. If imitation is a uniquely efficient and prepotent learning mechanism among humans, then rare inventions devised by an individual and spurious occurrences observed in nature can be 'seized upon' in a way that is not possible with more laborious and halting imitative abilities. Thus, technological advances can accumulate and rapidly spread through groups and across generations, as indeed they do within human cultures and today in our 'global village'. Note that with the vehicle of imitation, this dispersal effect is not dependent on deep causal (physical) understanding on anyone's part. By interpolation then, and without questioning their level of causal understanding, one could speculate that an advance in one or more of these special features of imitation was integral to hominins moving beyond the 'ape grade' and sustaining the seminal behaviour of stone knapping.

We would like to highlight yet another potentially-distinctive attribute of human imitative abilities that might have technological implications. This has to do with the fact that individuals learning through imitation do not limit themselves to acting with just the one object they saw demonstrated or to just the one purpose for which they saw it used. Instead, they generalize the learned behaviours to other similar objects and to analogous problem situations. People not only imitate but also adapt, improve, extend, and expand upon the goal-directed actions of others, or as the saying goes, they 'build better mouse-traps'. Even young children seem to abstract something more general from the imitative context and then generalize or 'transfer' it to other circumstances. This extension of activities learned through imitation may be part of what makes human tool-using so defining and productive, along with or instead of an advanced understanding of causality. Our current research on infant development focuses on this special transfer ability. In the following section, we describe several studies which document that young children not only attend to and imitate the means, as per Tomasello (1999), but also transfer the means to other novel objects and thereby 'invent' new goal-directed activities.

Research on transfer according to the means in human infants

Attention to the means during demonstrations

Our first approach to investigating transfer from instances of imitation was to simply look at what infants attend to while they observe demonstrations of object-directed actions (Brugger & Bushnell 1999). We showed 15- and 21-month-old infants mechanical toys which an adult model operated by manipulating a handle at one end. After seeing the demonstration, infants were given a turn to operate each toy themselves. The special feature of this study was that at least some versions of the toys were unusually long. For example, with the 'frog' toy, the handle at one end of a wooden base was some 80 cm away from a toy frog at the other end; the frog jumped back and forth when the handle was moved in and out. The extra length of the toys made it easy to see what the child looked at during the experimenter's demonstration, whether toward the experimenter's hands and their movements (i.e. the means) or toward the toy animal at the other end and its movements (i.e. the ends). Seeing how the model manipulated the handle would be necessary in order to 'imitate' rather than 'emulate' with these toys, following the distinction discussed earlier.

The results were that, first of all, the older children were generally more successful at making the toys work. As is shown in Table 20.1, this was true for each of the three toys used, that is, regardless of whether pressing, pulling, or turning was the critical action to perform. Overall, 21-month-olds reproduced the demonstrated effect by operating the handle on nearly 80 per cent of their trials, whereas 15-month-olds did so on only about one-third of theirs.

To further analyze these results, we broke the trials down according to whether the infant looked at the animal moving and also at the model's hands operating the handle during the demonstration or instead looked only at the animal moving. As might be expected, success was clearly predicated on having looked at the model's hands during the demonstration. Infants who looked at the model's hands were usually successful at working the toy themselves, while infants who did not look at the model's hands during the demonstration were almost never successful at working the toy themselves. This was true for the younger as well as for the older infants, suggesting that the motor demands of operating the handles were not too difficult for the 15-month-olds. Instead, the principal age difference was that older infants looked at the model's hands on almost every occasion, whereas younger infants failed to do so about 40 per cent of the time. This difference in attending to the means was the main reason why the younger infants were less successful at making the toys work overall.

Table 20.1. *Numbers of infants at each age who were successful or not at operating the toys.*

	Press		Pull		Turn	
	Dog toy		Frog toy		Cow toy	
	Yes	No	Yes	No	Yes	No
15-month olds	14	24	11	16	7	26
21-month olds	17	5	16	7	16	2

These results reinforce the distinction that Tomasello (1999) and Want & Harris (2002) have made between imitating the means and emulating the ends, and furthermore, they suggest that this distinction is relevant developmentally as well as between species. Only infants who looked at the model's hands were successful at working the toys themselves, and older infants looked at the model's hands more consistently than younger infants did. Looking to the hands by older infants may reflect an emerging understanding of physical causality or alternatively of social agency; either way, there is a cognitive 'burden' on the imitative learner— the infant has to know what to look for in a demonstration in order to benefit from it. It would be interesting to observe the gaze patterns of non-human primates during demonstrations with analogous 'long' objects where the means is performed at some distance from the ends. Perhaps the capacity to attend to the means despite the attraction of a perceptually salient ends originated with early hominins and was integral to their tool-making and tool-using skills.

Transfer according to the means with similar objects

In a subsequent study, we further examined the role of attention to the means during imitative learning and also began to explicitly address the question of transferring imitated behaviour beyond the initial learning situation (Sidman *et al.* 2001). In this work, we used a standard transfer design which in effect pitted attention to the means against attention to the ends. This allowed us to examine which aspect of an imitated behaviour — the means or the ends — might govern infants' later behaviour with a new toy.

The participants were 14- to 16-month-old infants, and each baby was first shown how to work two novel mechanical toys. As in the previous study, each toy had a handle on one end which could be manipulated in a certain way to move an animal mounted elsewhere on the toy. There were two kinds of handles (a faucet and a D-ring) and two ways that each of these could be manipulated (by pulling in and out and by turning in the fronto-parallel plane).There were also two different animals (a bird and a dog), and each of these could either jump up and down or spin around when the handle was manipulated. Each infant learned how to work a toy configured with one combination of these components and another toy configured with the alternative components. For example, an infant might first be shown how to *turn* the *D-ring* to make the *dog jump*, and then next shown how to *pull* the *faucet* to make the *bird spin*. For these learning trials, an experimenter demonstrated how to work the toy, and then offered the infant turns until he or she had successfully made the animal move.

After the infant had learned to work the two demonstration toys, a third object was brought out and offered to the infant with no demonstration. This test object was a hybrid of the first two; it had the handle of one demonstration object and the toy animal of the other. The question of interest was how the infant would try to work the test object, whether by turning or by pulling; that is, would they act on the test toy according to how they had acted before with that kind of handle (the means) or according to how they had acted before with that kind of animal (the ends).

The primary results are shown on Table 20.2. As can be seen, most infants acted on the test toy according to the means. If the test toy had the handle they had learned to pull earlier, they were more likely to pull than turn on the test trial, even though they had also learned earlier to move that particular animal by turning. These results suggest that similarity of the means, rather than similarity of the ends, governed motor transfer for these infants. Additional analyses showed that infants also tended to act on the test toy with the motor behaviour (pulling or turning) they were most proficient with during the learning phase and with the one they had used most recently.

Table 20.2. *Infants' performance according to means-driven or ends-driven behaviours.*

Action on test	Pull handle and turn toy	Turn handle and pull toy
Pulled	9	3
Turned	3	5

What is common to all of the strategies implied by the results is that the causal action itself is central, as opposed to the effect that is to be produced. Infants' transfer to the test toy was governed by the part of the object they had previously acted on, by the action they were most skilful at, and/or by the action they had most recently performed. The action was learned and practised in the context of imitating the experimenter with a particular toy, but infants seemed to abstract the action from this specific context and then apply it readily to another toy. As in the first study, attention to the means was evidently critical to infants' imitative behaviour, and in this case it also drove their subsequent interactions with a new toy which they had not seen demonstrated.

Transfer according to the means with dissimilar objects

We proceeded to investigate this transfer of learned means in another study (Bushnell & Sidman 2002), this time with more diverse objects and distinctive effects, in order to determine how far removed from the imitation context the means can be taken. In this study, we used two sets of three toys each. In one set, each toy could be activated by pressing a large blue button, and in the other, each toy could be activated by pulling on a white knob. Except that they had the same press-button or pull-knob, the toys in each set were otherwise different in colour, size, shape, and the nature of the effect they yielded — one toy made a song play, another made some balls spin, the third made a dog bark and wag its tail, and so forth.

The participants were again 14- to 16-month-old infants, and as before, each infant was first shown how to work two of the toys from a given (i.e. pressing or pulling) set. With each learning toy, an experimenter demonstrated how to work the toy, gave it to the baby for a turn, and then repeated this exchange once again. After these imitative turns with two of the toys, the third toy from the same action set was brought out and offered to the baby without any demonstration. The significant question was what the babies would do with the novel toy on this third or 'test' trial. Each baby had also been offered (with no demonstration) a toy from the alternative action set even before the learning trials began. These 'baseline' trials were included to see what infants would do with the stimulus toys spontaneously, in particular whether they could

Table 20.3. *Numbers of infants who successfully activated the toys in the baseline and test conditions.*

	Baseline	Test
Pull toys	3 (of 11)	8 (of 11)
Press toys	1 (of 11)	11 (of 11)
Total	**4 (of 22)**	**19 (of 22)**

discover the toys' cause–effect characteristics on their own, prior to demonstrations of any kind.

The results are shown in Table 20.3, which summarizes the infants' success at activating the toys on the baseline trials and on the corresponding test trials. As Table 20.3 shows, babies rarely managed to activate the toys on the baseline trials, although they typically did explore the toys and even touched the perceptually salient press-buttons and pull-knobs. However, on the test trials, almost all of the infants activated the toys successfully; furthermore, they often did so right away, within 3–4 seconds of being given the toy. That is, infants usually went straight for the test toy's button or knob and manipulated it in the same way (i.e. pressing or pulling) that had been effective with the demonstration toys. It is important to emphasize that this behaviour is not just the most obvious or easy thing to do with these toys, because on the baseline trials, babies almost never acted toward the same toys in this way. They clearly learned the effective action from imitating the experimenter with the first two toys, and then they just as clearly extended that action to the third toy which they had not seen demonstrated.

We had observed transfer of a learned means to a new toy in the previous study also; however, in the present case, the test toy was substantially different from either of the training toys. It was unlike them in colour, size, shape, the nature and number of discernable parts, and so forth, and there was nothing about it to suggest it would yield an effect like those of the training toys (which indeed it did not, it yielded a different effect). Only the manipulandum and the action required to operate it — that is, the means — were similar across the several toys; infants must have remembered this from the imitation trials and then carried it over to the new object.

Transferring the means: a proposed intermediate level of causal understanding

What we would like to emphasize from our work with infants is the apparent abstraction or disentangling of the pressing and pulling actions from the particular objects and effects for which they were first performed. From the imitation trials, infants seemed to learn not only how to play with those individual objects, but also that pressing buttons or pulling knobs is a generic 'way to make things happen'. Of course this does not reflect a real understanding of causality or any knowledge of how the toys worked in the intuitive physics sense. That is, we do not suppose that from the learning trials babies had come to understand the mechanics of springs or levers, so that when the novel toy came out, they could actually foresee what would happen if they pressed the button or pulled the knob. In most cases the mechanisms were hidden or actually electronic, so this kind of genuine insight would have been impossible for anyone, never mind for 15-month-olds. But at the same time, the infants seemed to have learned more than just that buttons are for pressing or knobs for pulling. Our impression is that on the test trials, they expected acting on the novel toy to have some sort of effect, although they did not (and could not) know exactly what it would be. It is almost as though in the infant's mind there was a blank slot to be filled in, in a kind of action–effect or means–ends frame.

In future research, we are planning to further investigate this frame-and-slot interpretation for early transfer. For example, we will observe the behaviour of infants when the test toy is 'broken', so that unlike the demonstration toys, pressing and pulling produces no effect and thus violates the presumed expectation for one. We will also observe infants' choice behaviour after demonstrations with one toy where an action leads to an effect and another where an action leads to no effect. But for now, if one accepts the frame-and-slot idea on speculation, an important implication from it is the opportunity for invention it affords.

As an analogy, consider the 'verb island' (Tomasello 1999) or 'pivot word' (Braine 1976) constructions that children produce at a certain stage in language acquisition. In these constructions, a certain key word is followed (or preceded) by any one of a whole class of words, as when a child uses the pivot term 'give' in combination with 'cookie', 'ball', 'hug', 'book' or any of the various content words he or she knows and might want an instance of. Such constructions are thought to reflect something less than genuine syntactic categories or sentence predicates, but they do grant the child an important measure of generativity which was not present before. In the prior phase, utterances consist of single word labels and so-called unanalyzed routines such as 'gimmeakiss' (but never 'gimme' anything else). These routines are highly context-dependent and impenetrable, and children's language at this point is somewhat imitative and robotic. However, when the pivot word strategy emerges, the child is suddenly a creative speaker, producing novel

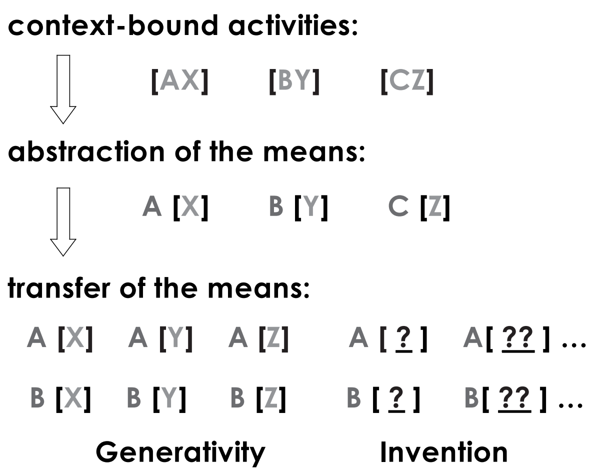

Figure 20.1. *Schematic depiction of the proposed developmental and perhaps evolutionary sequence, culminating in a generative action-grammar which enhances the discovery of new means–ends relationships.*

two-word utterances left and right and 'on the fly'. With this immature grammatical device, the child is in control of an infinitely expandable repertoire of expressible meanings.

We think that the transfer phenomenon we observed with our mechanical toys may represent the same kind of intermediate step but quantum leap with regard to children's means–ends activities with objects. Initially, effective actions and the objects to which they are applied may be inseparably bound together, essentially by conditioning processes, and the infant's goal-directed activities are limited to these learned responses and to the contexts in which they were learned. For example, Rovee-Collier and colleagues (Butler & Rovee-Collier 1989; Hayne *et al.* 1986; Rovee-Collier *et al.* 1985) have found that young infants easily learn a leg-kicking response which activates a mobile, but if even quite subtle aspects of the context or the mobile's elements are changed, infants do not generalize the kicking response to activate the new mobile. When an infant is able to abstract a means from its situational substrate, however, it becomes an entity in its own right and may be stored in an inventory of ways to make things happen. A particular action, say pressing or pulling, now plays the role of a pivot word. It is a sort of 'free radical' with the potential to combine with many objects, sometimes to no good effect at all and sometimes to produce exciting new effects. Thus because of this free-wheeling transfer, the odds of fortuitously discovering new means–ends relationships increase dramatically and further development is catalyzed. This hypothetical sequence and the generativity it permits are depicted in Figure 20.1.

The generalization process illustrated in Figure 20.1 is similar to developments originally described by Piaget (1952; 1954). As noted earlier, he wrote of infants first differentiating means from ends at about 8 months, and he characterized the subsequent 'stage' of cognitive development as 'the application of familiar schemata (actions) to new situations' (Piaget 1952, 212). Likewise, some have specifically defined 'play' as the free exploration of means with no particular ends in mind, and creativity and invention are supposed to arise from this aimless activity. Our contributions to this perspective are threefold. First, we have given a straightforward empirical demonstration of infants differentiating and transferring the means with our three-toy paradigm. Second, we emphasize that this process represents an intermediate stage in the development of causal cognition. When a learned action is applied to a new object, the child does not foresee or fully understand what effect will ensue, but nevertheless expects something to happen and performs the action with that ill-defined end in mind. Thus strictly speaking, transferring the means is not 'just play', with its explicitly purposeless quality. Finally, we emphasize the collaboration between adults and infants, between culture and cognition, that is involved with transferring the means. The actions which infants apply to new objects are not means they discover on their own, as Piaget might have it. They are low-probability behaviours which infants initially learn as effective actions by imitating others. But with their extraction from the learning context and application to new objects, these imitated behaviours then become the roots of individual invention, ingenuity, and problem solving. In our view, transferring the means makes for a primitive but generative action-grammar, by which individual infants can exceed what their elders have showed them.

Further speculation: transferring the means and the evolution of technology

To complete our discussion, we would like to speculate even further now and imagine a similar intermediate stage in the evolution of cognition, likewise characterized by the capacity to transfer learned means to new contexts. This conjecture rests on the several other parallels between cognitive development and evolution discussed earlier, namely with regard to future-oriented thinking, understanding connections between means and ends, advances in imitation, and so forth. As on these parameters we supposed that like infants early hominins lay somewhere between the abilities of non-human primates and modern mature humans, so we suggest they might have been positioned with respect to their capacities for invention. We have identified transfer according to the means as a potential mid-point in this regard in our work on infant development. In this final section and in what must be taken as a thought experiment, we apply this concept hypothetically to the situation of early hominins and their technology.

As a start, we know that non-human primates have at their disposal a number of purposeful and effective behaviours for attaining certain desired ends; termite fishing and nut cracking come to mind as obvious examples. Early hominins no doubt engaged in similar behaviours, including perhaps the earliest forms of using or even knapping stones as tools. These ape and early hominin behaviours presumably originated in some rare confluence of events and pressures to survive, and they are (were) sustained by their adaptive value and some form of cultural transmission such as Tomasello's (1999) emulation or Byrne's (1999) statistical learning. We suppose these activities may be the equivalent of unanalyzed routines in early language, with the causal action and its effect inseparably locked together as a unitary phenomenon, as illustrated in the top portion of Figure 20.1. Thus as with young infants' leg-kicking to mobiles, the action is not applied in other contexts and similarly the effect is not approached by other means. This state of affairs would make for a perpetuating but relatively changeless technology, as indeed some have characterized that of today's non-human primates and also that of hominins during the earliest phases of their evolution.

Following the parallel with development, we next suppose that at some point in prehistory, hominins gained the capacity to separate causal actions from their situational substrates, as exhibited by our 15-month-old research participants. Perhaps this advance was affiliated with the emergence of mimetic abilities in general (Donald 1991); like differentiating the means, imitation itself requires some sort of accessible mental representation of bodily actions. However, just when this change took place is less important for our purposes than the fact that, as in development, it would have increased the potential for invention and technological progress enormously. As is illustrated in the lower portion of Figure 20.1, with transfer according to the means, familiar means may be recombined with other familiar ends, to generate alternative ways of accomplishing something, say in the face of altered circumstances. Thus a hominin capable of transferring the means might strike an insect instead of a nut with a stone, for instance, and thereby 'invent' a more comfortable way to eat biting or pinching bugs. Conversely, he might poke a stick into a cracked nut instead of into a termite mound, and thus discover how to glean all the nutritious bits out of the shell. Familiar means could also be applied to novel objects, leading to entirely new and unforeseen effects which then could be capitalized upon. A hominin might hit grain instead of nuts with a stone and thus invent grinding flour. He might hit a stick instead of nuts with a stone and invent tent stakes or markers. He might hit a dangerous animal or other hominin with a stone and invent weapons. He might poke a stick into a hide instead of into a termite mound and invent primitive fasteners or buttons. And so on. De Beaune (2004) has similarly outlined how transfer processes may have engendered new technological skills in human evolution and gives further examples founded in lithic analyses.

Our position is that the inventions pictured in these scenarios were not initially grounded in sophisticated means–ends analyses or detailed knowledge of physical laws such as modern-day humans possess. That level of cognition came later, with capacities that allow the recognition of specific spatial, temporal, and other constraints that define individual causal relations. Recall that in development, understanding the real 'why' of particular means–ends relations seems to come after infants have been purposefully attaining those ends with those means for some time. Instead, the innovations imagined here could have derived from a more generic and 'vague' intuition that actions which are effective in one context may also be effective in others. As with infants, we suppose that this frame-and-slot level of understanding for cause and effect prompted the hominins in question to 'try' actions learned for one purpose on other novel materials and to attend closely to what the anticipated but unspecified effect turned out to be. Occasionally it turned out to be useful, and if the action sequence was then repeated and ultimately imitated by others, a technological advance had occurred, but it was not 'intentionally' (with foresight) invented. That is, the familiar means was intentionally tried in the novel context, but the exact outcome was not predicted beforehand, only realized afterwards and from then on, deliberately repeated.

The advantage of this intermediate level of causal cognition over say trial-and-error as a source of new skills has mainly to do with the 'trials' — they are more frequent, more systematic, and applied to a wider assortment of novel materials. Thus, stone knapping itself could have originated through transfer according to the means, perhaps through a series of 'accidents' involving percussive activities such as Marchant & McGrew (this volume) describe. Later, transfer could have been involved in the refinement of knapping techniques, such as in the use of bone and antler as hammers, for preparing edges, and so forth. With transfer according to the means, progress either to the initial incidence of stone knapping or subsequently from it did not have to await low-probability and fortuitous confluences of events; these were sought out and created in a more proactive way. Accordingly, the trajectory for technological change ramped up and ultimately produced the off-scale levels of technology enjoyed by even premodern humans in comparison to other primate species.

We might finally suppose that this process of transferring the means, of 'trying' effective actions out on novel objects and materials, was most intensely engaged in by the young in any given hominin group. They were after all not preoccupied with applying those actions to already known ends in order to survive, and yet they surely observed the elders doing so. The young could have imitated the elders' activities, and in their spare time may also have transferred the causal actions to new materials and thereby occasionally discovered useful new combinations. They would then have used these improved methods routinely as they grew to become adults, and the next generation in turn would imitate and expand upon them. Out of the hands of babes.

References

Baillargeon, R., 1994. How do infants learn about the physical world? *Current Directions in Psychological Science* 3, 133–40.

Bard, K.A. & J. Vauclair, 1984. The communicative context of object manipulation in ape and human adult–infant pairs. *Journal of Human Evolution* 13, 181–90.

Bates, E., V. Carlson-Luden & I. Bretherton, 1980. Perceptual aspects of tool-using in infancy. *Infant Behavior and Development* 3, 127–40.

Bayley, N., 1969. *Manual for the Bayley Scales of Infant Development*. New York (NY): The Psychological Corporation.

Boesch, C. & H. Boesch, 1983. Optimization of nut-cracking with natural hammers by wild chimpanzees. *Behaviour* 83, 265–86.

Boudreau, J.P. & E.W. Bushnell, 2000. Spilling thoughts: configuring attentional resources in infants' goal-directed actions. *Infant Behavior and Development* 23, 543–66.

Braine, M., 1976. *Children's First Word Combinations.* (Monographs of the Society for Research in Child Development 41(1).) Oxford: Blackwell.

Brugger, A.E. & E.W. Bushnell, 1999. Imitative Strategies Employed by 15- and 21-month-old Infants for Learning to Work Novel Objects. Poster presented at the Meetings of the Society for Research in Child Development, Albuquerque, New Mexico, April 1999.

Bushnell, E.W. & J.P. Boudreau, 1991. The development of haptic perception during infancy, in *The Psychology of Touch*, eds. M.A. Heller & W. Schiff. Hillsdale (NJ): L. Erlbaum, 139–61.

Bushnell, E.W. & J.P. Boudreau, 1998. Exploring and exploiting objects with the hands during infancy, in *The Psychobiology of the Hand*, ed. K.J. Connolly. (Clinics in Developmental Medicine 147.) London: MacKeith Press, distributed by Cambridge University Press, 144–61.

Bushnell, E.W. & J. Sidman, 2002. Beyond the Information Given: Extensions of Activities Learned Through Imitation. Symposium contribution presented at the International Society for Infant Studies, Toronto, Canada, April 2002.

Butler, J. & C. Rovee-Collier, 1989. Contextual gating of memory retrieval. *Developmental Psychobiology* 22(6), 533–52.

Byrne, R.W., 1995. *The Thinking Ape: Evolutionary Origins of Intelligence.* Oxford: Oxford University Press.

Byrne, R.W., 1999. Imitation without intentionality: using string parsing to copy the organization of behaviour. *Animal Cognition* 2(2), 63–72.

Byrne, R.W. & J.M. Byrne, 1993. Complex leaf-gathering skills of mountain gorillas (*Gorilla g. beringei*): variability and standardization. *American Journal of Primatology* 31(4), 241–61.

Canfield, R.L. & M.M. Haith, 1991. Young infants' visual expectations for symmetric and asymmetric stimulus sequences. *Developmental Psychology* 27(2), 198–208.

Chen, Z. & R.S. Siegler, 2000. *Across the Great Divide: Bridging the Gap Between Understanding of Toddlers' and Older Children's Thinking*. (Monographs of the Society for Research in Child Development 65(2), v-96.) Oxford: Blackwell Publishers.

Corballis, M.C., 1999. Phylogeny from apes to humans, in *The Descent of Mind: Psychological Perspectives on Hominid Evolution*, eds. M.C. Corballis & S.E.G. Lea. Oxford & New York (NY): Oxford University Press, 40–70.

Davidson, I. & W. Noble, 1993. Tools and language in human evolution, in *Tools, Language, and Cognition in Human Evolution*, eds. K.R. Gibson & T. Ingold. Cambridge: Cambridge University Press, 363–88.

de Beaune, S.A., 2004. The invention of technology: prehistory and cognition. *Current Anthropology* 45(2), 139–62.

Donald, M., 1991. *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Cambridge (MA): Harvard University Press.

Hayne, H., C. Greco, L. Earley, P. Griesler & C. Rovee-Collier, 1986. Ontogeny of early event memory: II. Encoding and retrieval by 2- and 3-month-olds. *Infant Behavior and Development* 9, 461–72.

Lederman, S.J. & R.L. Klatzky, 1987. Hand movements: a window into haptic object recognition. *Cognitive Psychology* 19, 342–68.

Leslie, A.M., 1984. Infant perception of a manual pick-up event. *British Journal of Developmental Psychology* 2, 19–32.

Lockman, J.J., 2000. A perception–action perspective on tool use development. *Child Development* 71(1), 137–44.

Matsuzawa, T., 1994. Field experiments on use of stone tools by chimpanzees in the wild, in *Chimpanzee Cultures*, eds. R.W. Wrangham & W.C. McGrew. Cambridge (MA): Harvard University Press, 351–70.

McCarty, M.E., R.K. Clifton & R.R. Collard, 1999. Problem solving in infancy: the emergence of an action plan. *Developmental Psychology* 35, 1091–101.

Meltzoff, A.N., 1988. Infant imitation after a 1-week delay: long-term memory for novel acts and multiple stimuli. *Developmental Psychology* 24(4), 470–76.

Meltzoff, A.N., 1995. Understanding the intentions of others: re-enactment of intended acts by 18-month-old children. *Developmental Psychology* 31(5), 838–50.

Meltzoff, A.N. & M.K. Moore, 1989. Imitation in newborn infants: exploring the range of gestures imitated and the underlying mechanisms. *Developmental Psychology* 25(6), 954–62.

Meltzoff, A.N. & M.K. Moore, 1994. Imitation, memory, and the representation of persons. *Infant Behavior & Development* 17(1), 83–99.

Meltzoff, A.N. & M.K. Moore, 1999. Persons and representation: why infant imitation is important for theories of human development, in *Imitation in Infancy*, eds. J. Nadel & G. Butterworth. Cambridge: Cambridge University Press, 9–35.

Mithen, S., 1996. *The Prehistory of the Mind: the Cognitive Origins of Art, Religion, and Science*. London: Thames & Hudson Ltd.

Nagell, K., R.S. Olguin & M. Tomasello, 1993. Processes of social learning in the tool use of chimpanzees (*Pan troglodytes*) and human children (*Homo sapiens*). *Journal of Comparative Psychology* 107(2), 174–86.

Piaget, J., 1952. *The Origins of Intelligence in Childhood*. New York (NY): International Universities Press.

Piaget, J., 1954. *The Construction of Reality in the Child*. New York (NY): Basic Books.

Rovee-Collier, C., P.C. Griesler & L.A. Earley, 1985. Contextual determinants of retrieval in three-month-old infants. *Learning & Motivation* 16(2), 139–57.

Schick, K.D. & N. Toth, 1993. *Making Silent Stones Speak*. New York (NY): Simon and Schuster.

Schlesinger, M. & J. Langer, 1999. Infants' developing expectations of possible and impossible tool-use events between ages 8 and 12 months. *Developmental Science* 2(2), 195–205.

Sidman, J., L.A. Lariviere & E.W. Bushnell, 2001. Factors Governing Infants' Transfer of Object-directed Actions. Poster presented at the meetings of the Society for Research in Child Development, Minneapolis, April, 2001.

Tomasello, M., 1999. *The Cultural Origins of Human Cognition*. Cambridge (MA): Harvard University Press.

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

Tomasello, M., E.S. Savage-Rumbaugh & A.C. Kruger, 1993. Imitative learning of actions on objects by children, chimpanzees, and enculturated chimpanzees. *Child Development* 64(6), 1688–705.

Toth, N., 1985. The Oldowan reassessed: a close look at early stone artefacts. *Journal of Archaeological Science* 12, 101–20.

Uzgiris, I.C. & J. Hunt, 1975. *Assessment in Infancy: Ordinal Scales of Psychological Development*. Champaign (IL): University of Illinois Press.

van Leeuwen, L., A.W. Smitsman & C. van Leeuwen, 1994. Affordances, perceptual complextity, and the development of tool-use. *Journal of Experimental Psychology: Human Perception and Performance* 20, 174–91.

Visalberghi, E., 1993. Tool use in a south american monkey species: an overview of the characteristics and limits of tool use in *Cebus apella*, in *The Use of Tools by Human and Non-human Primates*, eds. A. Berthelet & J. Chavaillon. New York (NY): Oxford University Press, 118–31.

Want, S.C. & P.L. Harris, 2002. How do children ape? Applying concepts from the study of non-human primates to the developmental study of 'imitation' in children. *Developmental Science* 5, 1–41.

Willatts, P., 1999. Development of means-end behaviour in young infants: pulling a support to retrieve a distant object. *Developmental Psychology* 35, 651–67.

Woodward, A., 1998. Infants selectively encode the goal object of an actor's reach. *Cognition* 69, 1–34.

Woodward, A. & J. Sommerville, 2002. Infants' Developing Sensitivity to the Intentional Structure of Action. Symposium contribution presented at the International Conference on Infant Studies, Toronto, Canada, April 2002.

Wynn, T., & W.C. McGrew, 1989. An ape's view of the Oldowan. *Man* 24, 383–98.