

Running head: DECISION MAKING IN INFANTS

5

Goal-Directedness and Decision Making in Infants

Ben Kenward, Sara Folke, Jacob Holmberg and Alexandra Johansson

Uppsala University

10

Gustaf Gredebäck

University of Oslo

Author Note

Ben Kenward, Department of Psychology, Uppsala University; Sara Folke, Department
15 of Psychology, Uppsala University; Jacob Holmberg, Department of Psychology, Uppsala
University; Alexandra Johansson, Department of Psychology, Uppsala University; Gustaf
Gredebäck, Department of Psychology, University of Oslo.

This work was supported by the European grant CONTACT (FP6-NEST 5010). We
thank everyone at Uppsala Infant Lab for much stimulating discussion and Claes von Hofsten
20 and Terje Falck-Ytter for comments on the manuscript.

Correspondence concerning this article should be addressed to Ben Kenward,
Institutionen för psykologi, Box 1225, 751 42 Uppsala, Sweden. E-mail:
ben.kenward@wolfson.oxon.org

Abstract

The term “goal-directed” conventionally refers to either of two separate process types – motor processes organising action oriented towards physical targets, and decision making processes which select these targets, by integrating desire for and knowledge of action outcomes. Even newborns are goal-directed in the first sense, but the status of infants as decision makers (the focus here) is unknown. In this study, 24-month-olds learnt to retrieve an object from a box by pressing a button, and then the object’s value was increased. After the object’s subsequent disappearance, these children were more likely to press the button to try to retrieve the object than were control 24-month-olds who had learnt to retrieve the object but for whom object value was unchanged. Such sensitivity to outcome value when selecting actions is a hall-mark of decision making. Fourteen- and 19-month-olds, however, showed no such sensitivity. Possible explanations include: they had not learnt the specifics of the action outcome; they had not acquired the necessary desire; they had acquired both but did not integrate them to make a decision.

Keywords: motivation; goal-directed action; decision making; infants; young children

Goal-Directedness and Decision Making in Infants

A great deal is known about the cognitive abilities which infants rely on when they carry out actions, and their capabilities in perception (Cohen & Cashon, 2003; von Hofsten, 2004), representation (Spelke, 2000), and memory (Rovee-Collier, Hayne, & Colombo, 2001) can be regarded as impressive. Less is known about the processes which motivate action selection in infants, because this is in some ways a more challenging problem. Consider, for example, a 6-month-old infant who immediately kicks and shakes a mobile which is once again tethered to its foot, 2 weeks after an initial training session (Rovee-Collier, 1999). While it is self-evident that the infant has perceived and remembered the mobile, there are numerous hypotheses for why this should motivate the infant to act (Rovee-Collier et al., 2001).

Two distinct processes for action selection, with separable neurological bases, exist in animals (Balleine & Dickinson, 1998). One process, termed goal-directed (Dickinson & Balleine, 2000; Hommel, 2003) motivates action by the integration of (a) an expectation that a specific action will have a specific outcome and (b) a desire for that outcome. Stimulus-driven *habitual action*, on the other hand, occurs as an automatic response to stimuli with which the action has become associated, for example through reinforcement learning (Balleine & Dickinson, 1998). Although this latter principle is simple, combinations of learnt responses can result in complex behaviour patterns (Donahoe, Burgos, & Palmer, 1993).

Note that this definition of goal-directedness, which focuses on the belief/desire processes which motivate action, is not the only definition with currency in psychology. The term *goal-directed* is often applied in a motor context to action which is organised with respect to a target object, manipulation, or posture (e.g. von Hofsten, 2004; Jeannerod, 2006). For clarity, we therefore refer to the belief/desire based process as a *decision making* process which *decides* upon action. To illustrate the difference, consider an infant attempting to grasp an object. The action is obviously goal-directed and may require sophisticated abilities such as object representation and trajectory prediction if the object is obscured or moving (Jonsson & von Hofsten, 2003). However, it is not as clear on the surface as to what has motivated the target selection. Decision making processes initiate goal-directed action by selecting targets, but so do other processes, such as habitual responses.

The tendency to decide upon a particular action depends on the current value of the outcome. Once a habitual action is learnt, on the other hand, it occurs in response to the stimulus,

irrespective of the current outcome value. This difference can be exploited to differentiate between the processes, using the experimental paradigm known as outcome revaluation.¹ When rats have learnt to obtain a particular food by pressing a lever, they will normally press the lever less after they have learnt that that particular food is now poisoned, indicating that their action is decided upon (Adams & Dickinson, 1981). If, however, the rats have received very many training sessions (Dickinson, Balleine, Watt, Gonzalez, & Boakes, 1995) or have lesions to the prefrontal cortex (Balleine & Dickinson, 1998), then the lever pressing continues irrespective of the reduced value of the outcome, which indicates that the action is a stimulus driven habit. As such, rats have been demonstrated to possess both decision making and habitual action selection mechanisms. In such experiments, the test is conducted in extinction (without presentation of outcome) which ensures that only the outcome representation learned during training can impact on performance. If the outcome is presented during the test then differences in performance may be due to reinforcement learning during the test phase.

The existence of decision making in adult humans requires no formal demonstration. The fact that stimulus-driven action is also commonplace (Bargh & Chartrand, 1999; Eimer, 1995; Hommel, 2000) perhaps requires discussion. It is common to perform a habitual action even though an unusual circumstance has rendered the action inappropriate (Reason, 1990). Pathologies in which such errors are enhanced are revealing. Adults with brain damage in areas such as the frontal lobes can be very prone to inappropriately perform habitual actions elicited by stimuli, such as involuntarily answering someone else's telephone despite instructions to the contrary (Humphreys & Riddoch, 2003; Rudd et al., 1998; Scepkowski & Cronin-Golomb, 2003). Such patients provide a clear demonstration of how processes governing goal-directed action can operate independently of decision making processes. Children have immature prefrontal cortices (Bell & Fox, 1992; Diamond, 2002; Sowell, Delis, Stiles, & Jernigan, 2001), indicating that, like these patients, their actions may be the products of decision making to a lesser degree than normal adults, and there are many cases in which they perseveratively perform

1. In fact the paradigm is often known as outcome devaluation, because outcome values are usually decreased not increased, but the more general term is appropriate because the logic of the method works both ways.

actions which are no longer appropriate (Hauser, 2003; Simpson & Riggs, 2007). For more discussion of how both decisions and non-decision-based processes motivate action in adults, see Jeannerod (2006), Hassin, Uleman, and Bargh (2005), and Maassen, Prinz, and Roth (2003).

It is widely accepted that infants' actions are goal-directed from an early age (Biro & Hommel, 2007). Most of the evidence cited in support of such claims demonstrates goal-directedness as we have defined it here in terms of motor organisation (e.g. Rochat & Striano, 2000; von Hofsten, 2004). Because the same term can refer to either motor organisation or decision making processes, however, and the intended sense is often not explicit, claims of goal-directedness may sometimes be open to over-interpretation. Without clear definition, evidence for goal-directed action may be interpreted as evidence for decision making. In general, however, a goal-directed action does not inform as to the presence of either decision making or habit, because both these processes must enlist goal-directed processes to organise action in pursuit of the target (Jeannerod, 1997). One aim of this manuscript is therefore to provide a rationale and framework for increased clarity in this issue.

Infant studies investigating means-ends behaviour, which might be classed as a type of decision making, do not usually rule out the possibility that the studied actions were motivated by stimulus-driven habits learnt through reinforcement. For example, the ability of infants to learn to pull a cloth to retrieve an out of reach object lying on it (Willatts, 1999) could be explained as follows. During initial manipulation of the cloth, the object is brought within reach. Because the object is desirable, this outcome is a reinforcer which serves to create an association between the stimulus of an object on a cloth and the action of pulling the cloth, so that in future the action will be elicited automatically by the perception of the stimulus.

We will describe two further cases of infants' complex goal-directed behaviour to illustrate that it need not involve decision making. The first example is when infants late in their second year reach around a transparent barrier using a two-step motion in order to grasp an object – a task which younger infants fail because they attempt to reach directly through the barrier for the visible object (Diamond, 1990). The second example is that 10-month-olds reach more slowly to grasp a ball when they will subsequently insert it into a tube than when they will subsequently throw it into a tub – inserting here requires more precision than throwing and is therefore facilitated by a slower preceding movement (Claxton, Keen, & McCarty, 2003).

What these tasks and many others have in common is that action is organised in a sequence using a forward model which anticipates and directs the positions of body parts relative to each other and to objects (see also Diamond, 2006; McCarty, Clifton, & Collard, 1999; von Hofsten, 2007). The later action steps are successful because they were anticipated when the first
 5 action steps were performed. The object behind the barrier can be grasped because the initial hand movement is not directly towards the object, but around the barrier. The ball can easily be inserted into the tube because the hand is not moving too fast following the reach. But importantly, these observations inform nothing as to what motivated the selection of the final goal object or action. In adults, such anticipatory goal-directed action is by default organised
 10 unconsciously once the target is selected (Jeannerod 2006) and as we saw earlier is separable neurologically and sometimes behaviourally from decision making processes.

In these two experiments, there are several potential non-decision-based motivations for reaching to grasp. Infants may have learnt a habit of reaching in response to graspable objects, because grasping is reinforcing. Alternatively, infants may inherit an automatic tendency to reach
 15 towards nearby objects (such a tendency would be adaptive for obvious reasons). When the reach for the ball precedes a second action of inserting or throwing, it may be that a reach has been automatically triggered by the circumstance where an action is to be performed with an object which is not yet grasped. The motivation for the final goal of inserting or throwing is also unknown. In this case the infants were imitating a demonstrator, and it is known that observation
 20 of an action can promote the automatic performance of the same action (Chartrand & Bargh, 1999; Dijksterhuis & Bargh, 2001; see also Craighero, Bello, Fadiga, & Rizzolatti, 2002).

To date only one study has attempted to demonstrate decision making in infants by using the outcome revaluation paradigm. Klosssek, Russell, and Dickinson (2008) allowed children to interact with a computer touch screen on which two icons were displayed. Touching one icon
 25 resulted in the display of a clip from a children's television program; touching the other icon resulted in a clip from a different program. After children were given the opportunity to learn these contingencies, they were habituated to the display of one or the other of the programs (by repeated and non-contingent presentation), in order to reduce its value as a reward. When subsequently tested for icon touching in extinction (i.e. without presentation of rewards),
 30 children older than 27 months preferred to touch the icon which prompted the non-devalued

reward, demonstrating that they expected and desired the non-devalued outcome: i.e. their actions were the result of decision making.

No such effect was found in younger children aged between 18 and 26 months. Thus, these children did not show the sensitivity to outcome value which is the key indicator of decision making. One interpretation is that they acted because of a habitual response to the display of the icon stimuli, learnt through reinforcement (Klossek et al., 2008), and although there are other possible explanations for this negative result, many of these were controlled for with additional experiments (for example, the habituation process did certainly reduce the reward value). Nevertheless, it is still possible that decision making emerges much earlier than 26 months. If this is the case then there must have been specific reasons why the youngest children could not be found to make decisions in Klossek et al.'s (2008) study, and it should be possible to discover an influence of outcome value on infants' action performance using a different method.

One aspect of Klossek et al.'s (2008) design was the necessity for participants to simultaneously consider whether to perform two different actions with different relative outcome values. This might have been a handicap for the younger children – it is known that children's ability to choose appropriate actions can be undermined by the introduction of similar but alternative possible actions (Frye & Zelazo, 2003). Here, we present a study of the effects of outcome revaluation on children's action performance which does not suffer from this potential drawback, and which in fact differs from Klossek et al.'s (2008) study in every substantial detail, including nature of task, nature of outcome, and method of outcome revaluation.

In a training phase, children learn to obtain an object from a box by pressing a button. In a subsequent play phase children learn how to use an object in combination with a play apparatus. This is intended to increase the value of the object because it is necessary to play the game. The play phase is terminated by the disappearance of the play object, and in a subsequent test phase the children's tendency to press the button is measured.

There are two experimental groups. For the revaluation group, the object is the same in the training and play phases, meaning that in the play phase the children learn that the object previously retrieved from the box is necessary for the game – this is intended to increase the value of the object in the box. For the control group, the objects are different in the two phases, meaning that the value of the object in the box remains at baseline. The hypothesis that children

make decisions predicts that the revaluation group will have a greater tendency than the control group to press the button in the test phase, because the outcome of the action (obtaining the object in the box) has become more desirable.

Experiment 1

5 *Method*

Participants

Participants were a self-selected sample who responded to an invitation letter sent to all families living in a Swedish city with children of appropriate age; therefore participants were mostly ethnically Swedish, exposed mostly to Swedish in the home, and had mixed
 10 socioeconomic backgrounds. The final sample consisted of 32 children of mean age 24.1 months ($SD = 8.0$ days), randomly assigned to the revaluation group (6 girls and 10 boys) and the control group (8 girls and 8 boys). Additionally 8 children were tested but excluded because of experimenter error or equipment failure (5 children), fussiness (2 children) or because the child never learnt to press the button (1 child).

15 *Materials*

Two different kinds of small object were used, counterbalanced across conditions: a white rubber ball (diameter 3.5 cm) and a plastic block covered with purple and orange tape (7 by 5 by 2 cm). During the training phase these objects could be obtained from a black button box (40 by 25 by 8 cm) with a button (QED Pal Pad, 15 by 12 cm) and a hatch which automatically
 20 swung open to reveal the object (the hatch was manually closed). The button was pressure sensitive with no moving parts because pilot studies showed that buttons with travel or noise are intrinsically rewarding to push. The button box was interfaced to a PC via a Phidgets 8/8/8 interface.

For each of the small objects there was an apparatus used in a game intended to increase
 25 the value of the object: a wooden ball run (110 by 70 by 15 cm) down which the ball could be rolled, bouncing back and forth, and a grey music box (50 by 35 by 20 cm), activated magnetically by placing the block on a marked target area on top. As long as the block was in place, a short repeated loop of children's music played and multicoloured lights flashed.

Procedure

30 Two experimenters carried out the procedure. Experimenter 1 carried out the training phase, and Experimenter 2 carried out the play and test phases (throughout the study they were

the same individuals). This rendered Experimenter 2 blind to which object was in the box in the training phase, and therefore unable to influence the child's behaviour in the test differently between conditions.

5 After greeting the child and the parent, the experimenters explained the purpose of the study, and the parent signed an informed consent form. The parent, child, and experimenters then entered the experimental room and began an informal warm-up phase intended to familiarise the child to the room and the experimenters. None of the objects used in the experiment itself were available in the warm-up phase. The experimental area (3.5 by 2.5 m) contained one chair, a video camera mounted on a shelf, the button box, either the ball run or the music box, a play mat, 10 and extra toys for the warm-up phase. The rest of the room (containing such things as the computer controlling the apparatus) was hidden behind a screen. When the experimenters judged that the child was sufficiently relaxed (usually after approximately 10 minutes of warm-up) Experimenter 2 left the room taking the warm-up toys, and the training phase was begun.

The child was seated on the mat facing Experimenter 1, with the button box between 15 them. The experimenter demonstrated that pressing the button caused the box to open and reveal the object, which she took and showed to the child. Then she put the object back and manually closed the box. The experimenter and the child then took turns opening the box and taking and replacing the object, with active participation of the experimenter varying according to the level of encouragement required by the child. When the experimenter was satisfied that the child was 20 capable of opening the box, the number of pushes required to open the box was increased, first to three pushes, and then later to five pushes once the child had learnt to open with three pushes. Each time the number of required pushes was increased, the experimenter demonstrated the new number of necessary pushes. The use of this ratio schedule was intended to maximise responding in the test phase.

25 The training phase ended when the child was capable of pressing the button five times to open the box. Occasionally the experimenter judged that the child was starting to lose motivation to participate, even though the child had not yet opened the box with five presses, in which case the phase was also ended. At the end of the phase the experimenter replaced the object in the box, told the parent to hold the child away from the apparatus, and left the room before 30 Experimenter 2 entered to begin the play phase.

On entering, Experimenter 2 showed the child a ball or block, demonstrated how it could be used with the appropriate play apparatus, and encouraged the child to join in the play. For the revaluation group, the object was the same type as had been in the box in the training phase (the original object itself remained in the button box). For the control group, the object was of the other type. When the experimenter judged that the child had reached a peak level of interest in the play, he made the object disappear by pretending to eat it (pilot tests had shown that this was the most effective way of removing an object without causing the child to search for it). This marked the beginning of the test phase, in which pushes on the button box (which was still on the mat) were recorded. The experimenter spent the remainder of the test session expressing regret for having eaten the object, and a desire to continue playing with the apparatus, if only a similar object could be found. Any surprise the child showed at the object being eaten was typically short-lived as the child became occupied by the search for a replacement.

The experimenter did not look at or draw attention to the button box until either the child began pressing the button, or until 60 s had elapsed without a press, at which point the experimenter began to ask if an object might be found in the button box. The button box could not in fact be opened during the test phase. If press rate is to be analysed, testing in extinction is necessary as it means test phase responses must be due to past experience rather than to reinforcement during the test phase. If no pressing occurred, the phase ended after 120 s; if pressing occurred, the phase ended 45 s after a press had last been registered. At the end of the experiment the object was informally made available to the child, so the procedure did not end on a note of frustration.

Throughout the procedure, parents were asked to sit passively on the mat or chair. They were allowed to encourage their child (verbally or by demonstrating) to take part in the training or revaluation phases if it was necessary to secure the child's participation. They were informed, however, that when the test phase began they were to be seated in a corner facing the wall (a text for them to read was on the wall to maintain their attention). This was to minimise their ability to influence their child's behaviour during the test phase.

Note that revaluation and control children are presented with one and two objects respectively in the training and revaluation phases, corresponding to a difference in environmental stimulation which in theory could effect activity level and therefore button pressing (Chipperfield & Eaton, 1992). Any such effect, however, could only be negligible,

because the difference is tiny compared to absolute stimulation levels during the child's visit to the lab (during the warm up phase, for example, children played with a collection of over 20 toys).

Analysis

5 Children who never pressed the button during training or never used the play apparatus were excluded from analysis. The first dependent variable was the latency to press the button in the test phase. To allow numerical rather than less powerful rank-based statistical analysis, children who did not press were allocated the maximum possible latency of 120 s. The second dependent variable was the test phase pressing rate, expressed as a percentage of the training phase press rate in order to reduce noise due to individual variation in baseline press rate.

10 The four possible combinations of objects used for the training and revaluation phases gave two revaluation sub-groups (ball then ball, block then block) and two control sub-groups (ball then block, block then ball). We included condition (revaluation or control) and training object (ball or block) as factors predicting the dependent variables in a balanced ANOVA. All statistical tests throughout this study are two-tailed.

We use a non-parametric permutation based ANOVA, the so-called DO test (Manly, 1997; Manly & Francis, 1999), in which the standard F -statistic is compared against a null-hypothesis distribution generated by random permutation (here, with 10,000 samples) rather than one generated by parametric assumptions. This was because the dependent variables did not satisfy such assumptions and could not be satisfactorily transformed. In such cases permutation tests are often more robust and more powerful than parametric and non-parametric alternatives (Good, 2004). For our variables the DO test was usually more powerful than parametric ANOVA, and we confirmed by simulation that it was no less robust against type I error (the error rate was estimated by repeatedly testing the data after transformation to make the null-hypothesis true, Peres-Neto & Olden, 2001).

Results and discussion

Because Experimenter 1 was not blind to the condition, it is possible that she might have unconsciously influenced the control and revaluation groups differently: in fact condition affected neither length of training phase, $F(1, 30) = 1.05, p = 0.350, \eta^2 = 0.03$, or rate of the child's training pushes, $F(1, 30) = 1.89, p = 0.180, \eta^2 = 0.06$ (see Table 1 for a training phase description). In the test phase, children in the revaluation group had a lower latency to start

pressing, $F(1, 28) = 8.8, p = 0.005, \eta^2 = 0.24$ (Figure 1, Table 2), and had a greater rate of pressing, $F(1, 28) = 5.6, p = 0.010, \eta^2 = 0.16$ (Figure 2, Table 2). There were no significant effects of whether the ball or the block was the training object, or any significant interaction effects. Particularly notable is that not a single child in the control group pressed the button in the test phase before being directed towards the box by the experimenter (figure 1), while 7 out of 16 in the revaluation group did so, $p = 0.007$ (Fisher's exact test). The stronger effect on latency than on press rate is probably because of rapid extinction in both groups – 81% of the children pressed only three times or less, and the median interval between the first and last press was only 1.0 s.

10 **Table 1**

Mean (SD) of Variables Describing the Training Phase

Age (months)	Length (s)	Push rate (per s)		Box openings triggered by child under ratio schedule of n pushes to one opening		
		Demonstration ^a	Child	$n = 1$	$n = 3$ ^b	$n = 5$ ^b
Experiment 1						
24	284 (276)	0.08 (0.06)	0.23 (0.12)	9 (10)	6 (4)	5 (3)
Experiment 2						
19	212 (68)	0.01 (0.01)	0.16 (0.09)	33 (18)	-	-
14	339 (175)	0.02 (0.01)	0.10 (0.06)	30 (12)	-	-

Note. $n = 32$ for all ages.

^aDemonstration pushes were primarily provided by the experimenter but occasionally by the parent. ^bThe ratio schedule was never increased beyond 1:1 in Experiment 2.

15 **Table 2**

Mean (SD) of Button Pressing Variables in the Test Phase

Age (months)	Latency to first press (s) ^a		Rate of pressing (%) ^b	
	Revaluation group	Control group	Revaluation group	Control group
Experiment 1				
24	68 (10.8)	105 (5.0)	35 (11.0)	7 (2.5)
Experiment 2				
19	81 (11.2)	84 (10.1)	7 (3.8)	22 (10.7)
14	81 (10.5)	90 (9.9)	28 (10.8)	38 (16.1)

Note. $n = 16$ for all groups.

^aIndividuals who did not press are allocated the maximum possible latency of 120 s.

^bIndividual's test phase pressing rates are expressed as a percentage of their training phase pressing rates.




Figure 1. Latency of 24-month-old control and revaluation group individuals to press the button in the test phase. The dashed line at 60 s indicates the point at which the experimenter prompted the child to examine the button box. The revaluation and control groups differ, $p = 0.005$ (see text).

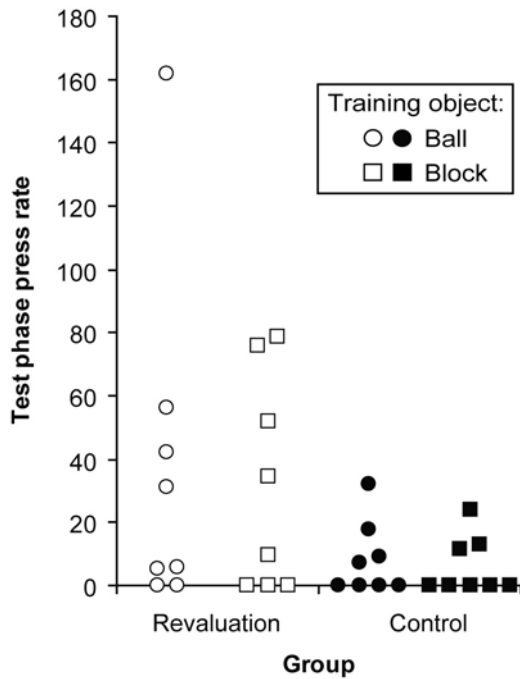


Figure 2. Button pressing rate of 24-month-old control and revaluation group individuals in the test phase, expressed as a percentage of each individual's training phase pressing rate. The revaluation and control groups differ, $p = 0.010$ (see text).

Therefore these 24-month-old children behaved as predicted if they are capable of decision making – there was a greater tendency to perform an action with a positive outcome, independent of the outcome value at the time the child learned the action. This is the first time such an effect has been found in children this young – Klossek et al. (2008) found an effect of

reevaluation on responding in a mixed age group in which the youngest children were 27 months. The next step is to determine if the same result could also be obtained in younger children.

Experiment 2

Method

5 *Participants*

Children were recruited in the same way as in Experiment 1. The final sample consisted of 32 children of mean age 14.0 months ($SD = 8.2$ days) randomly assigned to the reevaluation group (9 girls and 7 boys) and the control group (5 girls and 11 boys), and 32 children of mean age 19.2 months ($SD = 14.2$ days) randomly assigned to the reevaluation group (6 girls and 10 boys) and the control group (7 girls and 9 boys). Additionally 12 children (6 from each age group) were tested but excluded because of fussiness (2 children), experimenter error or equipment failure (2 children), a lack of interest in the play apparatus (4 children) or because the child never learnt to press the button (4 children).

Procedure

15 The procedure was identical to Experiment 1 with three modifications. An object preference phase was added to directly assess the effectiveness of the reevaluation process, in the following way. At the conclusion of the test phase, Experimenter 2 left the room and Experimenter 1 returned, presenting both a ball and a block to the child by placing the objects equidistantly from the child on the floor (with object side balanced over subjects) and asking
20 which of the objects the child would like to play with. The experimenter then observed passively for 60 s, except that if the child took neither object then the experimenter gave additional encouragement to play, and if after approximately 30 seconds the child had only played with one object, she asked if the child would also like to play with the other object.

Afterwards, the experimenter made a subjective judgement on a seven point bipolar scale
25 (1 = ball, 4 = neutral, 7 = block) about which object she felt the child preferred to play with, basing the judgement primarily on the time spent playing and the degree of positive affect displayed. Her judgement of this variable did not take into account her judgement of which object had probably been revalued, and these two judgements could differ. For example, if a child briefly rolled the ball in the ball run, but then played with the block for longer and with
30 more positive affect, then the ball had probably been revalued but the child was anyway judged to prefer the block. In order for this judgement not to be biased by her knowledge of the

condition, it was necessary for Experimenter 1 to be blind, necessitating a second procedure modification: throughout the experiment, both the ball run and the music box were in the room, rather than just one as before.

5 The third modification was the elimination of the ratio schedule, i.e., throughout the training phase, one push was sufficient to open the box. This change was made after pilot trials revealed that the younger infants often had trouble adapting to the ratio schedule although they easily learned to press the button once.

Analysis

10 The test phase was analysed as before, with the addition of age as a factor in the tests. A Manly and Francis (1999) DO ANOVA was used to determine the effects of age and revaluation object on Experimenter 1's judgement of object preference. Fisher's one sample tests with 10,000 randomised samples (Manly, 1997) were then used post-hoc to test for differences of preference from neutral, with alpha adjustment for multiple post-hoc comparisons (Benjamini & Hochberg, 1995).

15 *Results and discussion*

The training phase took longer for the 14-month-olds than for the 19-month-olds, $F(1, 62) = 14.52, p < 0.001, \eta^2 = 0.19$, and the 19-month-olds pressed at a greater rate, $F(1, 62) = 9.89, p = 0.002, \eta^2 = 0.14$ (Table 1). There were no significant effects of age, training object, condition, or any interactions on latency to push or the percentage push rate in the test phase
20 (Figures 3 and 4, Table 2). The trend sizes are negligible – the η^2 of condition is less than 0.004 for latency to push and 0.056 for percentage push rate (meaning condition explains just 0.4% and 5.6% of the variance respectively). The strongest trend (the higher push rate of the control group than the revaluation group in 19-month-olds) was in the opposite direction to that predicted if these children were making a decision to push to obtain the object.

25 We consider and discard the possibility that any of the three procedure modifications made since Experiment 1 could have been responsible for the lack of difference between revaluation and control groups at these ages. The preference phase was added after the test phase so could not have affected it. The presence of both play apparatuses rather than one, and the removal of the ratio schedule, could possibly have had some effect on press rates, but there is no
30 conceivable reason why these modifications would have acted to counteract a difference between the revaluation and control groups.

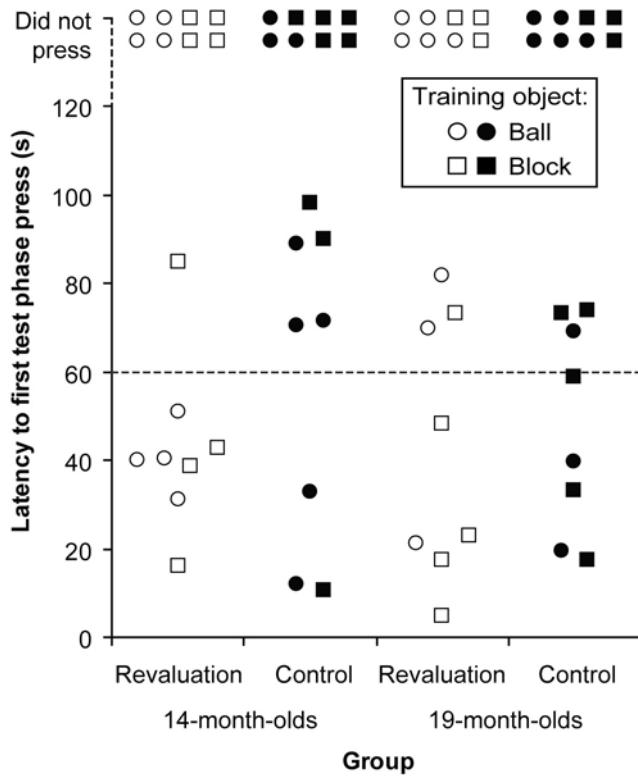


Figure 3. Latency of 14- and 19-month-old control and revaluation group individuals to press the button in the test phase. The dashed line at 60 s indicates the point at which the experimenter prompted the child to examine the button box. Neither age nor condition predicts latency (see text).




Figure 4. Button pressing rate of 14- and 19-month-old control and revaluation group individuals in the test phase, expressed as a percentage of each individual's training phase pressing rate. Some symbols representing a rate of zero are staggered downwards, for clarity. Neither age nor condition predicts rate (see text).

The absence of an effect of condition on pushing behaviour means there is no evidence that the 14- or 19-month-olds integrated an expectation of the outcome of button pushing with a desire for the outcome in order to make a decision. Although this could be because they are not able to make such a decision, there are numerous other potential reasons why they did not show an effect of outcome revaluation.

Firstly, the fact that some control group children had a low latency to push the button (25% of the 14- and 19-month-old control children pressed in the first 60 s of the test phase, compared with none of the 24-month-old control group) prompts the suggestion that children at this age were not capable of pushing the button any quicker even if they were motivated to do so. It is also the case, however, that half the control group children at both 14 and 19 months never pressed the button at all in the test phase, although they had all learnt to press in the training phase. In the light of this observation, this “floor-effect” explanation seems unlikely, although it cannot be ruled out.

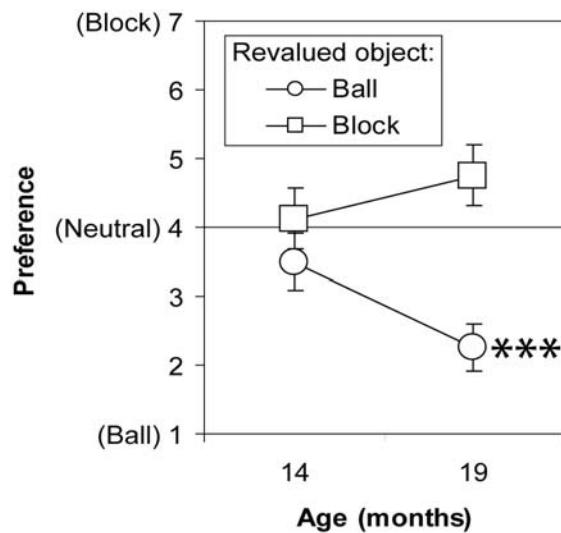


Figure 5. Mean object preference in the preference phase, as judged by Experimenter 1, by age and revalued object. $N = 16$ for each point. Error bars show standard error. Asterisks indicate a significant difference from neutral ($p < 0.001$).

Secondly, it is possible that the revaluation procedure was not effective at instilling an increased desire for the object in the box. The results from the object preference phase speak partially against this interpretation. The revalued object was a significant predictor of the child’s object preference, $F(1, 60) = 14.14$, $p < 0.001$, $\eta^2 = 0.18$, as was the interaction between age and revalued object, $F(1, 60) = 5.09$, $p = 0.028$, $\eta^2 = 0.06$, but not age itself, $F(1, 60) = 0.56$, $p = 0.454$, $\eta^2 = 0.01$. For each age group and revalued object, the trend is for the children to prefer the object which was revalued (Figure 5), although the preference was only significantly different from neutral for 19-months olds with the ball revalued, $p = 0.001$, $d = 0.54$. The fact

that revaluing the ball but not the block produced a significant difference from neutral preference may be due to the fact that there was an overall nonsignificant trend to prefer the ball, mean = 3.6, $p = 0.140$, $d = 0.35$.

Thirdly, it is possible that although the children learned to press the button, they had not learned to expect a particular object as a result of pressing. More results from the object preference phase speak partially against this interpretation. During this phase, the children would sometimes spontaneously take one of the objects, and place it on the button box (it could not be placed inside, because the button box was locked closed). This accords with the fact that during the training phase, the infants were encouraged not only to take the object from the button box but also put it back. The object taken was usually the “correct” object (i.e. the object which had been in the button box in the training phase). The 14-month-olds took the block three times correctly, and the ball once correctly and once incorrectly. The 19-month-olds took the correct object every time (eight balls, one block). The correct object was therefore taken significantly often when one considers all this data, $p = 0.002$, and also when one considers only 19-month-olds trained with the ball, $p = 0.008$ (sign tests).

These observations demonstrate that at least for 19-month-olds trained with the ball and in the revaluation group, there was an increased attractiveness of the ball, and an association between the ball and the button box. Because we only had good reason to believe that the training and revaluation may have been effective with these individuals, we repeated the analysis to determine the effect of condition on test phase pressing behaviour, using only data from 19-month-olds trained with the ball. There is still no effect of condition on latency to push, $F(1, 14) = 0.08$, $p = 0.755$, $\eta^2 = 0.00$, or pressing rate, $F(1, 14) = 2.6$, $p = 0.132$, $\eta^2 = 0.16$. In contrast, when data from the 24-month-olds trained only with the ball is reanalysed, there is still a significant effect of condition on latency to push,² $F(1, 14) = 5.0$, $p = 0.042$, $\eta^2 = 0.26$, although the effect on pressing rate is a non-significant trend, $F(1, 14) = 2.3$, $p = 0.095$, $\eta^2 = 0.14$.

2. In fact the paradigm is often known as outcome devaluation, because outcome values are usually decreased not increased, but the more general term is appropriate because the logic of the method works both ways.

It is unlikely that, although not detected due to lack of power, condition affects push latency in ball-trained 19-month-olds in the same way as in ball-trained 24-month-olds. Confidence distributions for the mean difference in latency due to revaluation were constructed for the ball-trained children in each age group using a randomisation procedure (Manly, 2007, p.18). Comparison of these distributions indicates a 97% confidence that a difference at least as great as that found in the 24-month-olds is not present in the 19-month-olds.

General discussion

In this study, 24-month-olds were first trained to push a button which opened a box revealing an object. After training, a revaluation group took part in a play session which increased the value of the object, and a control group took part in a play session which did not alter the object's value. Children subsequently pressed the button more quickly if the object's value had been increased. This observation is consistent with the hypothesis that these children made a decision to push the button because they believed it would have the result of obtaining the object. Younger children aged 14 and 19 months who took part in the same procedure did not show the same pattern – their latency to press the button was unaffected by the play session intended to increase the object's value. Further tests suggested that this was not because they had not learned that the object came from the button box, or because the revaluation procedure was ineffective. Neither of these controls was watertight, however, and there remain other possible explanations for the negative result which are not specifically based on a deficiency in decision making (see below).

Surprising though it may be that 24-month-olds but not 19-month-olds showed evidence of decision making in this rather simple task, these results accord closely (despite marked methodological differences) with the results of the only other study to use the outcome revaluation paradigm with children (Klossek et al., 2008). This lends extra validity to both studies. Our study detected sensitivity to outcome revaluation three month earlier, but this difference is possibly an artefact of the different age-groupings used. There is therefore little support for our suggestion that the necessity for the children to simultaneously consider two action-outcome pairings may have seriously impacted on children's ability to make decisions in Klossek et al.'s (2008) study.

If these results do reveal a developmental transition in decision making ability, then this would parallel much of what is known about children's abilities to solve tasks involving the

application of abstract rules or geometrical planning. For example, children cannot consistently rotate objects in order to fit them into correspondingly shaped holes until around 22 months of age (Örnkloo & von Hofsten, 2007). In the delayed non-matching to sample task, in which the participant retrieves an object under a cover with a symbol which is different to a sample symbol
5 which has previously been displayed, children do not succeed until roughly 21 months of age (Diamond, Towle, & Boyer, 1994). More generally, this period of the late second year is a time when many profound developmental milestones are reached, including the production of the first word combinations (Anisfeld, Rosenberg, Hoberman, & Gasparini, 1998), the establishment of pretend play (Lillard, 2002), and mirror self-recognition (Amsterdam, 1972; Nielsen &
10 Dissanayake, 2004).

Our experimental design was a simplified version of the standard outcome revaluation procedure, in that the participants only learnt one action in the training phase, and groups were compared in their absolute tendency to perform the action, rather than in a tendency to prefer one action over another. This design was necessary to rule out the possibility that negative results
15 could be accounted for by confusion due to multiple choices (Frye & Zelazo, 2003). However, it allows the theoretical possibility of an alternative hypothesis to explain the results. Given the prerequisites of an association between the button box and the object it contained, an increased attraction to the revalued object, and a learnt action of button pressing in response to the button box, children in the revaluation group might have been more attracted to the button box (because
20 of the association with the attractive object) and the button box might therefore have elicited more pushing (see Colwill & Motzkin, 1994). In a two-action design this hypothesis can be ruled out if both actions are performed on the same manipulandum, because then no single response can be habitually associated with it (Dickinson, Campos, Varga, & Balleine, 1996). When
Klossek et al. (2008) modified their procedure to include this technique, they still found an effect
25 of revaluation, although below 36-months the trend was not significant.

The 19-month-olds showed all the prerequisites for this hypothetical mechanism to operate, but because there were no group differences, there is no evidence for its operation. This implies that whichever mechanism is responsible for the group differences becomes functional late in the second year. Because this hypothetical alternative mechanism is very simple, it seems
30 unlikely that it should develop at this relatively late age. It is therefore reasonable to conclude that the most likely explanation for sensitivity to outcome devaluation in our 24-month-olds is

that they were making decisions. Klosssek et al. (2008) and Balleine and Ostlund (2007) discuss details of the cognitive and neural bases of decision making.

Motives for button pushing in the 14- and 19-month olds.

5 If the control children were already pushing the button as fast as they were capable, then an increased motivation to push could be present in the revaluation children but behaviourally silent. This seems unlikely, however, because half of the control children at these ages did not press at all in the test phase. Assuming it is not the case, it can be concluded from the absence of group pressing differences that the revaluation children cannot have decided to press based on a belief that pressing would result in obtaining the object, and a desire for the object greater than
10 that of the control children.

In any case, children of all ages may have been motivated to press by decision making processes based on desires and beliefs different to those the experiment was designed to detect. Some control group 24-month-olds also pressed the button in the test phase, and while these actions may have been habitual, the conclusion that revaluation group 24-month-olds made
15 decisions implies that those in the control group may also have decided to press the button but for different reasons. So might the 50% of 14- and 19-month olds who pressed the button in the test phase.

Although the 19-month-olds did have an increased tendency to take the revalued object when presented with both objects, this does not prove that they desired it more than the other
20 object in the absence of either. Likewise, although they could associate the box with the object obtained from it, it is not proven that they remembered that pressing the button would retrieve that object. Possible decision based motives for their button pushing therefore include that they wanted to open the hatch, and that they either did not remember or were indifferent as to which object would be obtained. It is also conceivable that they had an incorrect belief about what was
25 in the box – in the absence of a clear memory they may have assumed it contained the object they were being encouraged to seek. Even in the absence of any knowledge about the consequences of pressing the button, these infants may have decided to press for exploratory reasons in the belief that pressing the button would allow them to (re-)learn the consequences of the action. The possibility that the younger children's different performance was due not to
30 deficiencies in decision making, but in memory, would be consistent with the observation that the younger control children's rates of test-phase pushing were high compared with the 24-

month-old control children: a less clear memory of the consequences could provide an increased incentive to button press for exploratory reasons.

As described in the introduction, however, behaviour which seems to have resulted from a decision, because it achieves a result which is presumably rewarding to the child, may in fact be motivated by non-decision-based processes. One such possibility (especially with the 14-month-olds) is that the infants were producing the stereotyped surface banging action which is a predictable consequence of normal maturation (Thelen, 1979), and can be seen as exploratory even in a non-decision-based sense because it allows contingency and affordance learning. This mechanism is very unlikely to be solely responsible because, although we did not formally measure the baseline rate of button pushing without demonstration, piloting indicated it would be close to zero.

A second possibility is that the infants were still imitating the actions of Experimenter 1 in the training session (this was presumably why the infants pushed during training, at least initially). Many types of and mechanisms for imitation exist, and they can be automatic and non-decision-based (Brass & Heyes, 2005; Dijksterhuis & Bargh, 2001). Note that even infants' abilities to appreciate constraints on others' actions (Gergely, Bekkering, & Kiraly, 2002) and to deduce the goal of an observed incomplete action (Meltzoff, 1995) need not involve decision making, because these abilities relate to how target actions are accomplished, not how they are selected.

Another plausible explanation is that these children were motivated by a learnt habitual response to the button stimulus. Infant studies often produce results which are challenging to explain in terms of decision making, but straightforward to explain as a response to a stimulus with stimulus generalisation (Ghirlanda & Enquist, 2003). For example, once 6- to 8-month-olds have learnt to pull a cloth and retrieve a toy resting on it, they will also pull the cloth when the toy is not on it, or pull the cloth without retrieving the toy, sometimes using the same pattern of actions as those classified as "intentional" in the toy-on-cloth context (Willatts, 1999).

Perseverative repetition of previously rewarded actions in a context in which they are no longer appropriate to achieve the original goal are common in infants in a variety of studied situations (Hauser, 2003), and also occur in older children (e.g. Simpson & Riggs, 2007). Some of these studies even use methods resembling outcome revaluation. For example, McCall and Clifton (1999) trained 8½-month-olds to open a box and retrieve an object. Testing them in the

dark, so that the object could not be seen even after the box was opened, they found that the infants were equally likely to open the box, and to reach inside for the object, whether or not they had seen that the box contained an object prior to lights-out a few seconds earlier.

Infants certainly can expect the outcome of their actions (Elsner & Aschersleben, 2003; Hauf, 2007; Kenward, 2008), so it is quite possible that, as the preference phase tests suggest but do not prove, the younger infants here did learn the consequence of button pressing (and did remember at the time of testing), but were nevertheless unable to utilise this knowledge in a decision making process. This would represent a *décalage*, in which information is possessed but is not used to guide action (Hauser, 2003), as seen for example in the A not B error in which infants reveal knowledge of object location through gaze patterns which they do not use to guide manual object search (Hofstadter & Reznick, 1996; Piaget, 1954).

The inadequacy of natural outcome revaluation experiments.

Observations which resemble a positive result in an outcome revaluation experiment can frequently be made in young children's day to day life. For example, children will begin to engage in behaviour which obtains food, such as pointing at the freezer and saying "bread", early in their second year. The tendency to perform this behaviour certainly depends on the value of the outcome – food is more valuable when one is hungry, and this is when it is more likely to be requested.

This observation, however, does not allow conclusions regarding decision making, because, unlike in a controlled revaluation experiment, the child has had plenty of experience of asking for and receiving bread under different levels of hunger. This experience corresponds to different strengths of reinforcement depending on the level of the internal hunger stimulus. It is therefore plausible that the food requesting behaviour could be learnt as a habitual response to a strong internal hunger stimulus. Note that this interpretation does not imply that the child does not desire bread, but that such a desire does not motivate a decision.

As discussed in the introduction, reaching in itself is not a demonstration of a decision. Nonetheless, it can further be observed that infants' tendency to reach may relate to the desirability of the object (e.g. Moses, Baldwin, Rosicky, & Tidball, 2001), and this situation again resembles an outcome revaluation procedure. Like hunger, however, attractiveness may function as an internal stimulus which the individual has learnt to respond to by reaching. The plausibility of this hypothesis presents a problem for attempts to establish infants as decision

makers, because while it may well be that infants decide to reach, for the reasons given here outcome-revaluation does not lend itself to the study of simple reaching. The introduction of an arbitrary contingency allows the outcome revaluation method to be applied in a valid fashion, but, as the data presented here can be interpreted as suggesting, such an arbitrary contingency may interfere with infants' ability to make decisions.

Final remarks.

We began by drawing the distinction between two very different types of process which have both been termed goal-directed. We chose to retain the term goal-directed to describe motor processes which organise action around physical targets. It has been argued, however, that the use of terms such as 'goal' in describing processes for which the motives are not well understood may be misleading (Sirois & Jackson, 2007). Nevertheless, the term is well established in the literature in the motor control context, as are other terms such as 'intention', 'plan', and 'expectation'. Nonetheless, it is true that these terms do have the potential to confuse because their everyday uses have connotations of decision making. We therefore advocate that it is always made explicit in which sense these terms are meant, in both motor and decision making contexts.

In this study we have shown that there is good reason to accept that 24-month-olds can make decisions, because their tendency to act depended on the expected outcome value even after outcome revaluation. The status of younger children as decision makers remains uncertain because a relationship between outcome value after revaluation and tendency to act could not be demonstrated, either here or in the only other study to apply outcome revaluation to children (Klossek et al., 2008). Processes such as automatic imitation or habitual learning can also therefore explain the action observed. Although these hypotheses provide simpler putative mechanisms than decision making for the behaviour observed in the younger children, the simplest mechanism is not necessarily the most parsimonious in a developmental context, in which complex processes may be operating undetectably due to their immaturity. A challenge, however, remains – those who would like to establish that infants make decisions must find a situation in which infants can pass the outcome revaluation test, or find a valid new assay for decision making.

References

- Adams, C. D., & Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *Quarterly Journal of Experimental Psychology*, *33B*, 109-121.
- Amsterdam, B. (1972). Mirror self-image reactions before age two. *Developmental Psychobiology*, *5*(4), 297-305.
- 5 Anisfeld, M., Rosenberg, E. S., Hoberman, M. J., & Gasparini, D. (1998). Lexical acceleration coincides with the onset of combinatorial speech. *First Language*, *18*(53), 165-184.
- Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: Contingency and incentive learning and their cortical substrates. *Neuropharmacology*, *37*(4-5), 407-419.
- 10 Balleine, B. W., & Ostlund, S. B. (2007). Still at the choice-point: Action selection and initiation in instrumental conditioning. *Annals of the New York Academy of Sciences*, *1104*, 147-171.
- Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, *54*(7), 462-479.
- 15 Bell, M. A., & Fox, N. A. (1992). The relations between frontal brain electrical-activity and cognitive-development during infancy. *Child Development*, *63*(5), 1142-1163.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate - a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B-Methodological*, *57*(1), 289-300.
- 20 Biro, S., & Hommel, B. (2007). Becoming an intentional agent: Introduction to the special issue. *Acta Psychologica*, *124*(1), 1-7.
- Brass, M., & Heyes, C. (2005). Imitation: Is cognitive neuroscience solving the correspondence problem? *Trends in Cognitive Sciences*, *9*(10), 489-495.
- Chartrand, T. L., & Bargh, J. A. (1999). The Chameleon effect: The perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, *76*(6), 893-910.
- 25 Chipperfield, J. G., & Eaton, W. O. (1992). Reactivity and Environmental Stimulation as Predictors of Motor-Activity Level in Children. *Personality and Individual Differences*, *13*(5), 591-601.
- Claxton, L. J., Keen, R., & McCarty, M. E. (2003). Evidence of motor planning in infant reaching behavior. *Psychological Science*, *14*(4), 354-356.
- 30

- Colwill, R. M., & Motzkin, D. K. (1994). Encoding of the unconditioned stimulus in Pavlovian conditioning. *Animal Learning & Behavior*, 22(4), 384-394.
- Craighero, L., Bello, A., Fadiga, L., & Rizzolatti, G. (2002). Hand action preparation influences the responses to hand pictures. *Neuropsychologia*, 40(5), 492-502.
- 5 Diamond, A. (1990). Developmental time course in human infants and infant monkeys, and the neural bases, of inhibitory control in reaching. *Annals of the New York Academy of Sciences*, 608, 637-676.
- Diamond, A., Towle, C., & Boyer, K. (1994). Young childrens performance on a task sensitive to the memory functions of the medial temporal-lobe in adults - the delayed-
 10 nonmatching-to-sample task - reveals problems that are due to non-memory-related task demands. *Behavioral Neuroscience*, 108(4), 659-680.
- Diamond, A. (2002). Normal development of prefrontal cortex from birth to young adulthood: Cognitive functions, anatomy, and biochemistry. In D. T. Stuss & R. T. Knight (Eds.), *Principles of frontal lobe function* (pp. 466-503). London, UK: Oxford University Press.
- 15 Diamond, A. (2006). The early development of executive functions. In E. Bialystok & F. Craik (Eds.), *Lifespan cognition: Mechanisms of change* (pp. 70-95). NY: Oxford University Press.
- Dickinson, A., & Balleine, B. W. (2000). Causal cognition and goal-directed action. In C. Heyes & L. Huber (Eds.), *The evolution of cognition* (pp. 185-204). Cambridge, MA: The MIT
 20 Press.
- Dickinson, A., Balleine, B. W., Watt, A., Gonzalez, F., & Boakes, R. A. (1995). Motivational control after extended instrumental training. *Animal Learning and Behavior*, 23, 197-206.
- Dickinson, A., Campos, J., Varga, Z. I., & Balleine, B. (1996). Bidirectional instrumental
 25 conditioning. *Quarterly Journal of Experimental Psychology Section B-Comparative and Physiological Psychology*, 49(4), 289-306.
- Dijksterhuis, A., & Bargh, J. A. (2001). The perception-behavior expressway: Automatic effects of social perception on social behavior. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 33, pp. 1-40). San Diego: Academic Press.
- 30 Donahoe, J. W., Burgos, J. E., & Palmer, D. C. (1993). A selectionist approach to reinforcement. *Journal of the Experimental Analysis of Behavior*, 60(1), 17-40.

- Eimer, M. (1995). Stimulus-response compatibility and automatic response activation: Evidence from psychophysiological studies. *Journal of Experimental Psychology: Human Perception and Performance*, 21(4), 837-854.
- 5 Elsner, B., & Aschersleben, G. (2003). Do I get what you get? Learning about the effects of self-performed and observed actions in infancy. *Consciousness and Cognition*, 12(4), 732-751.
- Frye, D., & Zelazo, P. D. (2003). The development of young children's action control and awareness. In J. Roessler & N. Eilan (Eds.), *Agency and self-awareness: Issues in philosophy and psychology* (pp. 244-262). Oxford: Oxford University Press.
- 10 Gergely, G., Bekkering, H., & Kiraly, I. (2002). Rational imitation in preverbal infants. *Nature*, 415(6873), 755-755.
- Ghirlanda, S., & Enquist, M. (2003). A century of generalization. *Animal Behaviour*, 66, 15-36.
- Good, P. (2004). *Permutation, parametric, and bootstrap tests of hypotheses* (3rd ed.). New York: Springer.
- 15 Hassin, R. R., Uleman, J. S., & Bargh, J. A. (Eds.). (2005). *The new unconscious*. Oxford: Oxford University Press.
- Hauf, P. (2007). Infants' perception and production of intentional actions. In *Progress in brain research: From action to cognition* (Vol. 164, pp. 285-301): Elsevier.
- Hauser, M. D. (2003). Knowing about knowing - dissociations between perception and action systems over evolution and during development. *Annals of the New York Academy of Sciences*, 1001(79), 79-103.
- 20 Hofstadter, M., & Reznick, J. S. (1996). Response modality affects human infant delayed-response performance. *Child Development*, 67(2), 646-658.
- Hommel, B. (2000). The prepared reflex: Automaticity and control in stimulus-response translation. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 247-273). Cambridge, MA: MIT Press.
- 25 Hommel, B. (2003). Acquisition and control of voluntary action. In S. Maasen, W. Prinz & G. Roth (Eds.), *Voluntary action: Brains, minds, and sociality* (pp. 34-48). Oxford: Oxford University Press.

- Humphreys, G. W., & Riddoch, M. J. (2003). Fractionating the intentional control of behaviour: A neuropsychological analysis. In J. Roesler & N. Eilan (Eds.), *Agency and self-awareness: Issues in philosophy and psychology*. Oxford: Oxford University Press.
- Jeannerod, M. (1997). *The cognitive neuroscience of action*. Oxford: Blackwell.
- 5 Jeannerod, M. (2006). *Motor cognition: What actions tell the self*. Oxford: Oxford University Press.
- Jonsson, B., & von Hofsten, C. (2003). Infants' ability to track and reach for temporarily occluded objects. *Developmental Science*, 6(1), 86-99.
- Kellman, P. J., & Arterberry, M. E. (1998). *The cradle of knowledge: The development of perception in infants*. Cambridge, MA: MIT Press.
- 10 Kenward, B. (2008). 10-month infants visually anticipate an outcome contingent on their own action. *Manuscript submitted for publication*.
- Klossek, U. M. H., Russell, J., & Dickinson, A. (2008). The control of instrumental action following outcome devaluation in young children aged between 1 and 4 years. *Journal of Experimental Psychology-General*, 137(1), 39-51.
- 15 Lillard, A. (2002). Pretend play and cognitive development. In U. Goswami (Ed.), *Blackwell Handbook of childhood cognitive development* (pp. 188-205). Oxford, UK: Blackwell.
- Maassen, S., Prinz, W., & Roth, J. (Eds.). (2003). *Voluntary action: Brains, minds, and sociality*. Oxford: Oxford University Press.
- 20 Manly, B. F. J. (1997). *Randomization, bootstrap and Monte Carlo methods in biology*. London: Chapman & Hall.
- Manly, B. F. J., & Francis, R. (1999). Analysis of variance by randomization when variances are unequal. *Australian & New Zealand Journal of Statistics*, 41(4), 411-429.
- McCall, D. D., & Clifton, R. K. (1999). Infants' means-end search for hidden objects in the absence of visual feedback. *Infant Behavior & Development*, 22(2), 179-195.
- 25 McCarty, M. E., Clifton, R. K., & Collard, R. R. (1999). Problem solving in infancy: The emergence of an action plan. *Developmental Psychology*, 35(4), 1091-1101.
- Meltzoff, A. N. (1995). Understanding the intentions of others - reenactment of intended acts by 18-month-old children. *Developmental Psychology*, 31(5), 838-850.

- Moses, L. J., Baldwin, D. A., Rosicky, J. G., & Tidball, G. (2001). Evidence for referential understanding in the emotions domain at twelve and eighteen months. *Child Development, 72*(3), 718-735.
- Nielsen, M., & Dissanayake, C. (2004). Pretend play, mirror self-recognition and imitation: A longitudinal investigation through the second year. *Infant Behavior & Development, 27*(3), 342-365.
- Örnkloo, H., & von Hofsten, C. (2007). Fitting objects into holes: On the development of spatial cognition skills. *Developmental Psychology, 43*(2), 404-416.
- Peres-Neto, P. R., & Olden, J. D. (2001). Assessing the robustness of randomization tests: examples from behavioural studies. *Animal Behaviour, 61*, 79-86.
- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.
- Reason, J. (1990). *Human error*. Cambridge: Cambridge University Press.
- Rochat, P., & Striano, T. (2000). Perceived self in infancy. *Infant Behavior & Development, 23*(3-4), 513-530.
- Rovee-Collier, C. (1999). The development of infant memory. *Current Directions in Psychological Science, 8*(3), 80-85.
- Rovee-Collier, C., Hayne, H., & Colombo, M. (2001). *The development of implicit and explicit memory*. Amsterdam: John Benjamins.
- Rudd, R., Maruff, P., MacCuspie-Moore, C., Velakoulis, D., Brewer, W., & Pantelis, C. (1998). Stimulus relevance in eliciting utilisation behaviour: Case study in a patient with a caudate lesion. *Cognitive Neuropsychiatry, 3*(4), 287-298.
- Scepkowski, L. A., & Cronin-Golomb, A. (2003). The alien hand: cases, categorizations, and anatomical correlates. *Behavioral and Cognitive Neuroscience Reviews, 2*(4), 261-277.
- Simpson, A., & Riggs, K. J. (2007). Under what conditions do young children have difficulty inhibiting manual actions? *Developmental Psychology, 43*(2), 417-428.
- Sirois, S., & Jackson, I. (2007). Social cognition in infancy: A critical review of research on higher order abilities. *European Journal of Developmental Psychology, 4*(1), 46-64.
- Sowell, E. R., Delis, D., Stiles, J., & Jernigan, T. L. (2001). Improved memory functioning and frontal lobe maturation between childhood and adolescence: A structural MRI study. *Journal of the International Neuropsychological Society, 7*(3), 312-322.
- Spelke, E. S. (2000). Core knowledge. *American Psychologist, 55*(11), 1233-1243.

- Thelen, E. (1979). Rhythmical stereotypies in normal human infants. *Animal Behaviour*, 27, 699-715.
- von Hofsten, C. (2004). An action perspective on motor development. *Trends in Cognitive Sciences*, 8, 266-272.
- 5 von Hofsten, C. (2007). Action in development. *Developmental Science*, 10(1), 54-60.
- Willatts, P. (1999). Development of means-end behavior in young infants: Pulling a support to retrieve a distant object. *Developmental Psychology*, 35(3), 651-667.