ELSEVIER

Cognition 67 (1998) 287–309

COGNITION

Children's mental representation of referential relations

I.A. Apperly, E.J. Robinson*

School of Psychology, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

Received 26 May 1997; accepted 20 April 1998

Abstract

We identify a surprising discrepancy in children's performance in two tasks which appear superficially to require handling of the same properties of the representational mind. Four- to six-year-olds made judgements about the knowledge of a protagonist who had only partial information about an object: the child knew that an object in a box had two descriptions, X and Y (e.g. dice and eraser), but the protagonist had access to only one of these, X. In Experiment 1, children who passed a standard false-belief task also judged correctly that the protagonist did not know the X was Y, but often judged wrongly that he did know there was a Y in the box. In Experiment 2, children predicted wrongly where the protagonist would look for a Y: the problem was not purely linguistic. We argue that success on standard theory-of-mind tasks can be supported by a more basic representing ability than is assumed in current theories, and that children's mental representation of referential relations between the world and the mind subsequently undergoes important change.

Keywords: False belief; Theory of mind; Child

1. Introduction

In contemporary discussions of folk psychology, the mind is normally treated as a representational system. There are several important consequences of this view which together capture some of the essence of what differentiates the mental world from the physical. First, the lack of a concrete causal link with the real world allows the mind to entertain non-existing states of affairs, giving us the

* Corresponding author. E-mail e.j.robinson@bham.ac.uk

0010-0277/98/$19.00
PII S0010-0277(98)00030-4

power of imagination. The price of such freedom is the capacity to misrepresent reality and thus be committed to the truth of a non-existent state of affairs. The third important consideration is that '...all representation – whether it is done by the mind, language, pictures, or anything else – is always under certain aspects and not others...' (Searle, 1983): when an object, for example, is 'represented', only a limited number of the items in the extensional set belonging to that object are captured.

Each of these consequences raises problems of reference when we want to think about (or otherwise represent) mental contents. Our theories of folk psychology must therefore also explain how these are handled, and developmental accounts must explain this acquisition: How do we come to represent the relation between the world, and contents that don't exist in 'current here-and-now reality' or else only capture some of the properties of their real world referents? Could all of these problems be solved at a stroke with the advent of a mature theory of the mind, or is there explanatory mileage in allowing that they may differ in their developmental complexity? On the basis of data from two investigations of related skills that are unusually late in developing, we shall argue that there is. We begin, though, by contrasting some developmental accounts of what is commonly known as a 'theory of mind'.

A classic demonstration of errors in children's handling of mentalistic problems is the unexpected transfer task (Wimmer and Perner, 1983). In the original version, the child observes as the location of a chocolate is changed in the absence of the protagonist, Maxi. On Maxi's return the child is asked to predict where Maxi will look for the chocolate. To an adult it is clear that Maxi will look in the old location. However, children rarely begin to pass this task until their fifth year, younger children reply that Maxi will look in the current location of the chocolate.

Consistent results on this and other related tasks (e.g. Gopnik and Astington, 1988) have led to the apparent changes at 3–4 years becoming the primary focus for investigation and theorising. Broadly speaking, explanations divide into two kinds. The first claims that the afore-mentioned tasks underestimate children's understanding (e.g. Fodor, 1992; Robinson and Mitchell, 1992, 1994, 1995; Leslie, 1994; Mitchell, 1994, 1996; Moore et al., 1995; Russell, 1996). Variously, they try to explain away the stage-like appearance of the change, emphasising instead, factors that may develop gradually across the period of interest. For example, Mitchell and Lacohee (1991) present evidence that children's problems actually consist in a reality bias; Russell et al. (1994) found evidence that children's errors were non-mentalistic and executive in nature; Fodor (1992) argues that children are merely conservative in selecting hypotheses for explaining behaviour. The second kind of explanation characterises children's development as a change in their concept of the mind (e.g. Chandler, 1988; Flavell et al., 1981; Wellman, 1990; Perner, 1991; Gopnik and Wellman, 1992; Carpendale and Chandler, 1996). It is this type of explanation that makes the strongest claims about young children's handling of referential relations, and thus forms the appropriate backdrop to our investigations. However, in the discussion we will also consider whether performance difficulties could explain our findings.

Perner claims that the transition at 3–4 years of age marks a conceptual leap in children's understanding of the mind and of representations in general. By meta-representing mental states as representations, the sophisticated referential relations between a representation and the world are captured (e.g. Perner, 1991). Theories such as Perner's impressively accommodate the co-emergence of a number of folk-psychological skills such as predicting behaviour on the basis of a false belief and understanding deception. However, a survey of published experiments in the broader area of mentalistic development gives the impression of a more complex picture than these theorists suggest.

Besides the well-known clustering of changes at around the turn of the 3rd year, there is a second cluster of tasks that children rarely pass before 6–7 years. For example, children up to the age of around 5 or 6 years tend to judge that ambiguous verbal input is more informative than it really is (e.g. Flavell et al., 1981; Robinson and Robinson, 1982; Robinson and Whittaker, 1987; Robinson, 1994; Carpendale and Chandler, 1996). They also appear to have related problems judging the informativeness of pictures (e.g. Robinson and Robinson, 1982; Chandler and Helm, 1984; Taylor, 1988). Finally, Russell (1987) presents data showing that children aged 6–7 have difficulty understanding intensional contexts in which the normal rules of extensional substitution do not hold.

The strategy of Perner and others (Perner and Davies, 1991) has been to try and demonstrate that if task-specific limiting factors are removed, children can indeed succeed on these apparently harder tasks at only 4 years of age (but see Mitchell et al., 1997; Robinson, 1994). However, by insisting that the unexpected transfer task diagnoses the last qualitative change in children's understanding of the mind they seem tied to this explanation for all such tasks. We will argue that this is an important weakness in such accounts.

An alternative tack taken by Chandler (e.g. Carpendale and Chandler, 1996), is to synthesise the results from interpretation and ambiguity tasks into a picture of late conceptual development at 6–7 years. Children's early abilities are not given the singular importance they have in Perner's account. Instead, Chandler argues that the pattern of success and failure at this age can be best explained if children posses a simple 'copy theory' of the mind (Chandler and Boyes, 1982). Later success on the second cluster of tasks corresponds to a qualitative change involving insight into 'the interpretive character of the knowing process'.

Amongst the experimental results clustering at 7 years, Russell's look at children's understanding of intensionality fits the least easily with Chandler's account (Russell, 1987). In Russell's study, children aged 5–7 years were told a number of different vignettes in which a character came to have incomplete knowledge about a particular situation, e.g. George has his watch stolen but doesn't see the thief. Some test questions asked what it was possible to say about the character's thoughts, e.g. 'can we say that George was thinking...?'. In such intensional contexts, the answer was constrained by the knowledge state of the character so that if he did not know that the thief was a man with curly red hair, it was incorrect to say that George was thinking 'I must find the man with curly red hair who stole my watch'. Results showed a significant age-related improvement, with errors on the intensional ques-

tions decreasing from an average of 87.5% at age 5, to around 50% at age 7. In contrast, virtually all the children were able to distinguish between intentional and extensional contexts. That is, in line with the expected performance of this age group on the standard false-belief tests, they were able to ignore their own knowledge of reality and answer that we cannot say that 'George knows that the thief is a man with curly red hair'.

Although it has the necessary time course, Chandler's account appears ill-equipped to cope with this dissociation. Children who can acknowledge George's ignorance of the thief's hair colour seem already to be beyond a simple copy theory of the mind. A simple copy theory would lead them to assume that if George knows something about the person who stole his watch, he must know everything. But if this is the case, their misunderstanding of the intensional context cannot be due to their possessing a copy theory. Perhaps an elaborated copy theory of some kind would allow them to deny that we can say that George knows the thief is a man with curly red hair, but then they should surely also deny that we can say that George knows that the man with the curly red hair stole his watch. Either way, like Perner whose 4-year-old's theory of representation should entail the restrictions on reference in intensional contexts, Chandler seems forced to try and explain away children's problems. So how might Russell's findings be explained?

The relative complexity of Russell's vignettes and questions leaves it open for both Perner and Chandler to object that incidental task factors could indeed be masking children's true competence. Although children's success at reporting that George doesn't know that the thief is a man with curly red hair weakens this argument somewhat, the differences between Russell's tasks and standard theory-of-mind tasks are too great for such a possibility to be ignored. A different strategy is to rule intensionality out of court as a mentalistic phenomenon.

According to Russell, 'understanding opaque [intensional] contexts is a form of specifically linguistic knowledge that is probably not present before middle childhood' (Russell, 1992; p. 496). This view contrasts with Perner's general account because handling of opacity is seen by Russell as a problem posed in language interpretation, not in mental representation. However, Perner himself uses just such an explanation to reconcile these data with his theory: By claiming that the problem Russell poses is a pragmatic one, he places it outside of the 'theory-of-representation' domain of his own account (Perner, 1991). A similar option is also open to Chandler.

An alternative interpretation becomes apparent when the particular referential problems posed by Russell's tasks are examined closely. The intensionality of contexts such as those in Russell's study results from the conditions of satisfaction of a meta-representation (how we describe George's belief) being dependent upon the features of the primary representation (the belief itself) rather than the features of the object of that representation (see Searle, 1983). That is to say, the problems of reference in intensional contexts arise because representations only partially capture their referents. Perner's over-arching theory does not differentiate the source of these referential problems from complexities that arise because the mind can refer to different times, places and hypothetical states of affairs. Yet there is no a priori

reason for treating these as equivalent in development: representing a mental state as only partially capturing its referents may not be equal in difficulty to representing its referring to different times, places or hypothetical situations. We believe that these distinctions are exactly what make some mentalistic tasks easier than others.

Unlike the above account, an explanation that splits children's achievements in this way might be extended to explain their problems with other tasks. For example, ambiguous utterances can occur because, as representations of the world, messages only capture certain aspects of their referents. Hence, in an ambiguous sentence, it is impossible to distinguish the intended referent from another similar object which meets the same limited specifications of the description. The current account thus suggests that ambiguity tasks are harder than false-belief tasks, because they require the child to model the contents of the ambiguous utterance as partial and therefore limited in their referential scope.

In sum, the respective theories of Perner and Chandler each appear to have two possible ways of accounting for Russell's findings: either children's problems with intensionality (referential opacity) are incidental and task-specific, or they are purely linguistic and thus come outside of either theory of mental development. If it can be shown that neither account holds, and if children's errors do not just reduce to performance problems, the way would seem to be clear for our alternative explanation. Therefore, in Experiment 1 we tested the strength of Russell's important results in a setting that bears much closer comparison with standard false-belief tests and therefore permits stronger conclusions to be drawn about differences in difficulty. In Experiment 2 we went on to examine Russell's argument that referential opacity is a purely linguistic notion (Russell, 1992).

2. Experiment 1

As summarised above, Russell's tasks required children to reflect on whether a particular linguistic description was true or false: can we say George was thinking...? (Russell, 1987). We aimed to create a simplified version of Russell's task that retained the essential features of intensional contexts arising out of the partial knowledge of one story character, while eliminating some task-specific complexity. We reduced the number of embedded clauses in the test question from four to one. In our modified versions, children were not asked to judge whether we can say such and such about George's thoughts, but were simply asked whether the protagonist knew such and such. Hence the form of our test questions was very similar to those which have been widely used in theory-of-mind research (e.g. Pillow, 1989; Pillow, 1993; Pratt and Bryant, 1990). We used objects referred to by two definite descriptions (X and Y). The child and experimenter knew both descriptions whilst the puppet protagonist was only able to find out one (X) when he looked at the object. Thus description X was the only correct content of the puppet's propositional attitudes towards the object. We assessed children's understanding of this with three questions. Following Russell, we considered whether children were willing to use description 'Y' in an extensional context: 'can Heinz see the 'Y' in the box?'.

We also checked that the children could differentiate intentional and extensional contexts by asking, 'does Heinz know the 'X' is 'Y'?' To create an intensional context we asked: 'does Heinz know there's a 'Y' in the box?' Recall that the children in Russell's experiments found intensional contexts particularly hard. By reducing task-specific demands we expected this to decrease somewhat. However, our foregoing analysis would predict that any questions regarding Heinz's knowledge should remain harder than the standard false-belief tasks, since both appear to concern the partial nature of mental representations. If this prediction was borne out, even with our modifications of Russell's tasks, then we would be in a strong position to argue that handling of referential opacity is not acquired in a single conceptual change as Perner's general theory holds (Perner, 1991).

2.1. Method

Each child was tested individually on four tasks. Two were like the task described above and it was not appropriate to judge that the protagonist knew the second description on the basis of his limited informational access. A third, control task was similar in form, but this time it was appropriate to judge that the protagonist knew the second description even though he had the same limited informational access. The fourth task was a standard deceptive box task.

2.1.1. Participants

We tested 28 children from a nursery class (13 boys and 15 girls) aged between 3 years 7 months and 4 years 9 months (mean age 4 years 3 months), and 48 children from a reception class (27 boys and 21 girls) aged between 4 years 10 months and 5 years 11 months (mean age 5 years 4 months). We shall refer to these as the 4- and 5-year-olds, respectively. All the children attended the same junior/infant school with a lower middle class/upper working class catchment area in Birmingham, UK, and spoke English as their first language.

2.1.2. Materials

A bouncy ball, a rubber dice ('rubber' is the standard British term for an eraser) and a toy duck were all contained in similar tin boxes. During the experiment, each item was referred to by two possible definite descriptions: 'ball'/'present', 'dice'/'rubber' and 'duck'/'toy' respectively. In the case of the first two items, only one description was obvious from visual inspection ('ball' and 'dice'). For the third, both 'duck' and 'toy' were discernible from visual access and this was used in the control task. The fourth item, used for the deceptive box task, was a plastic horse which was contained in a packaging box with a picture of a teddy on the outside. The protagonist for every trial was a puppet called Heinz.

2.1.3. Procedure

Trials with the ball and dice followed the same general form. Children were first allowed to look inside the box and the visually-obvious label was agreed (children had no difficulty in naming the ball and the dice). Next, the second label had to be

agreed. In the case of the ball, the children were conspiratorially informed that '...this is going to be a present for Heinz, except we haven't told him and we don't want him to find out right now, so we'll have to whisper very quietly when he looks...' For the dice, they were allowed to feel the item, and if they did not spontaneously identify it as a rubber, it was demonstrated until this second label was agreed. The two descriptions were then reiterated: '...so it looks like a dice and it feels like a rubber...' Children had therefore heard referential use of both labels for each item and the different modes of perceptual access had been made clear. Next the children observed as Heinz looked inside the box. The lack of other perceptual access was emphasised in the case of the ball by whispering the questions, and in the case of the dice by saying 'now Heinz is going to look inside the box but he's not going to feel'.

With the box still open, half of the children were asked:

Q1 (extensional) can Heinz see the (present/rubber) inside the box? (correct = yes).

The other half had no question, guarding against the possibility of this initial 'extensional' question promoting later errors. After the box was closed the children were asked either:

Q2 (intentional) does Heinz know the (ball/dice) is a (present/rubber)? (correct = no).

or

Q3 (intensional) does Heinz know there's a (present/rubber) in the box? (correct = no).

The trial with the duck/toy was similar. However, in this case only one label, 'duck', was agreed in advance, to control against children allowing substitution merely because the labels had been agreed with the experimenter. Thus, the 'toy' label was novel to the children when they were asked Q1: 'can Heinz see the toy inside the box?' (correct = yes). Next they were asked Q1a: 'can you see the toy with your eyes right now?' (correct = no, as the box was tipped away from the child). All children were asked Q3: 'does Heinz know there's a toy inside the box?' (correct = yes) followed by Q3a: 'do you know there's a toy inside the box?'(correct = yes). Thus, this trial served as a double control, ensuring firstly that the child was willing to allow substitution of definite descriptions when the context was transparent, and secondly that the child was able correctly to attribute the basic seeing–knowing relationship both to themselves and the puppet.

On the deceptive box trial children were shown the box and asked what they thought was inside. After agreeing upon 'teddy', the children were allowed to look inside the box and find the horse. After closing the box, children were asked 'Heinz hasn't seen inside this box before. When he first sees it, before he opens it, what will he think is inside?'

The four trials were partially counter-balanced: the deceptive box test always came first or last and the duck/toy control test always came before the dice and ball tasks. The order of the latter two tasks was counter-balanced and the question type (intentional or intensional) varied between participants. In sum, aside from simple differences in order, there were four distinct trial types for the ball and dice items. Trial type was consistent within child: intentional only, extensional and intentional, intensional only, and extensional and intensional.

2.2. Results and discussion

First we checked children's answers in the duck/toy control task. Two children were excluded from the 4-year-old sample after they insisted that they could see the toy when the box was tipped away from them (control question, Q1a). Two children were excluded from the 5-year-old sample for the same reason and a further three for evidencing suspicions about the puppet's ability to see or know anything. All other children got the four questions of the duck/toy control task correct.

Performance on the deceptive box task was in line with the published literature: ten of the 26 4-year-olds (38%) answered correctly, compared with 33 of the 43 5-year-olds (76%). As would be expected, performance on the deceptive box improved significantly with age $\chi^2(1,n = 69) = 10.11$, $P < 0.002$.

For the other two tasks, the dice/rubber and ball/present, there were no significant order effects (chi-squared test, all $P > 0.43$). Table 1 gives a break-down of children's performance in the four basic experimental groups. To test whether the presence or absence of the extensional question Q1 ('can Heinz see...?') had any effect upon responses to intentional or intensional questions, children were classified according to whether they got none, one or both of these (Q2 or Q3) questions correct. No effect was found (chi-squared test, all $P > 0.389$).

On the basis of this result, data from children who were asked Q1 were combined with those from children who were not. Table 2 shows the incidence of correct responses to the two questions in the dice/rubber and ball/present tasks.

Children performed at ceiling in answer to the extensional question: they all judged correctly that Heinz could see the present or rubber. This could be interpreted as a sign that children did not over-extend the consequences of the puppet's ignorance of those descriptions of the items to extensional contexts created by the verb 'see'. However this same result would be predicted if children were generally relaxed about referential substitution in any context, and as we shall see below their answers to the intentional and intensional questions suggest that this was the case.

Our main interest is in the answers to intentional Q2 (e.g. 'does Heinz know the dice is a rubber?') and intensional Q3 (e.g. 'does Heinz know there's a rubber in the box?'). As shown in Table 2, our younger children performed poorly in answer to

Table 1
Numbers of 4- and 5-year-olds who got none, one or both intentional (or intensional) questions correct in the four principle conditions of Experiment 1

Condition	Number of correct responses in nursery class (4 years) $n = 26$			Number of correct responses in reception class (5 years) $n = 43$		
	0	1	2	0	1	2
Intentional only	2	2	2	1	3	7
Intentional + extensional	5	1	1	3	2	5
Intensional only	4	1	1	7	2	2
Intensional + extensional	7	0	0	10	1	0

Table 2
Numbers (%) of 4- and 5-year-olds who answered intentional and intensional questions correctly in the dice/rubber and ball/present tasks in Experiment 1

Age group	Question	Dice/rubber	Ball/present
Nursery (4 years)	Q2, intentional 'knows X is Y?' $n = 13$	6 (46%)	3 (23%)
	Q3, intensional 'knows Y is in box?' $n = 13$	2 (16%)	1 (8%)
Reception (5 years)	Q2, intentional 'knows X is Y?' $n = 21$	14 (67%)	15 (71%)
	Q3, Intensional 'knows Y is in box?' $n = 22$	3 (14%)	4 (18%)

both questions, with no significant differences between questions. In contrast, the older children showed highly-significant differences between answers to intentional and intensional question types, with the intensional question ('does Heinz know there's a rubber/present in the box?') the more difficult. For the dice/rubber, χ^2 (1, $n = 43$) = 12.6, $P < 0.0001$ and for the ball/present, χ^2 (1, $n = 43$) = 12.35, $P < 0.0001$. Many children judged correctly that Heinz did not know that the dice was a rubber (or the ball a present) while few judged correctly that Heinz did not know there was a rubber (or a present) in the box.

These differences for the 5-year-olds were reflected in comparisons with the deceptive box test. There was no difference in difficulty between intentional Q2 and the deceptive box task, while intensional Q3 was harder than the deceptive box task: McNemar's χ^2 (1, $n = 22$) = 13.1, $P < 0.001$ for the dice/rubber, and χ^2(1, $n = 22$) = 12.1, $P < 0.001$ for the ball/present. These results are consistent with the assumption that intentional Q2, like the deceptive box test, tapped children's ability to distinguish between intentional contexts (as represented by the mind) and extensional contexts (the real state of affairs).

Comparing the age groups, although performance on the intentional Q2 ('knows X is Y?') showed signs of improvement from an average of 34% to 69%, this did not reach significance when the frequencies of responses (neither Q2 correct, one Q2 correct or both Q2s correct) were compared: χ^2 (2, $n = 34$) = 5.12, $P = 0.077$. Performance on the intensional Q3 ('knows Y is in the box?') did not improve significantly: average scores across the two Q3s were 3/26 (11%) to 7/43 (15%).

In line with the findings of O'Neill et al. (1992) and Robinson et al. (1997), our 5-year-olds were able to distinguish between the knowledge acquired from seeing versus feeling, although this appeared to be a little harder than the deceptive box task. Additionally, in the present study, they used this modality-specific access to infer the puppet's knowledge of alternative descriptions for the objects, as assessed by intentional Q2. However, the dramatic disparity with performance on intensional Q3 suggests that even though children could correctly identify what the puppet did or did not know, they often failed to apply this in intensional contexts where substitution of co-referential terms affected the truth value of the sentence.

To summarise, the surprising result from Experiment 1 is that children who are successful at predicting another's false belief, correctly report another's ignorance of a label for an object and are significantly worse at treating descriptions of what a person knows as constrained by that person's actual knowledge about an object. While performance on the false-belief task and the intentional Q2 showed improvement across our age groups, performance on the intensional task remained consistently bad. This poor performance in response to intensional Q3 replicates the essential finding of Russell (1987) even though our tasks were dramatically simplified versions of his.

According to Perner (1991), success on false-belief tests marks a qualitative shift in children's understanding of the representational mind, and acknowledgement of false belief is considered to entail the same skills as understanding referential opacity (intensionality). What, then, is the basis of this enormous discrepancy between children's ability to contrast explicitly a person's knowledge and ignorance of two definite descriptions of an object (X and Y), and their ability to acknowledge that it is wrong to make 'Y' the content of the person's propositional attitudes?

Maybe there was simply a performance difficulty with the intensional question, in that children for some reason failed to make use of the fact that Heinz did not know the dice was a rubber. That is, whereas the intentional Q2 asked children directly about Heinz's knowledge of the item's properties, perhaps the intensional Q3 asked them only indirectly to take Heinz's ignorance into account. Such an explanation for the difference in difficulty would predict an effect of presentation order on answers to the two questions, i.e. if children were asked the intentional question first they should perform better on the intensional question than if they were asked the intensional question first. In Experiment 1 we were concerned not to confound any such effect with overall differences between the two types of question. Having found such differences, Experiment 2 was designed to allow the above prediction to be tested.

Both Perner (1991) and Russell (1996) have suggested a language-based explanation for why children might answer 'yes' to an intensional question while answering 'no' to an intentional one. Whilst the latter unambiguously concerns Heinz's knowledge vis à vis definite descriptions of the object (e.g. that the dice is a rubber), the former could be heard as 'does Heinz know there's an (object that you and I know is, and therefore can refer to as, 'a present') in the box?' which has the correct answer 'yes'. As Russell noted in his 1987 article, the distinction between transparent and opaque interpretations is a purely pragmatic one, so in certain circumstances a transparent interpretation could be the more appropriate (Russell, 1987). Perhaps children have problems deciding which interpretation to make.

For such an explanation to hold water we need to ensure that adults do indeed observe the anticipated convention in our scenarios. Only then could children's responses be considered non-conventional. We therefore presented 33 sixth-form students (17–18 years) with the dice/rubber and ball/present tasks from Experiment 1. We were only interested in responses to the intensional question but as in Experiment 1, for half of the sample this was preceded by an extensional question. We anticipated that the relative familiarity of dice/rubbers would allow some adults to

say that one could tell just by looking that it was a rubber, and this was borne out in the results: for the dice/rubber only 25/33 (76%) said that Heinz did not know there was a rubber inside the box. For the ball/present there were no familiarity problems and this time 30/33 (91%) said correctly that he did not know there was a present inside the box. Thus, overall we were happy that the adult interpretation of our scenarios was as anticipated, which allows the explanation in terms of linguistic pragmatic difficulties to remain a possibility.

Yet another possible reason for children's errors was suggested in the introduction: the representational demands of intensional contexts may be qualitatively different from those of the theory-of-mind tasks passed at around 4 years because they require representation of the fact that mental representations are partial. Initially, this suggestion appears inadequate since the intentional question above (e.g. 'does Heinz know that the ball is a present?') surely asks about partial representation but was no harder than the deceptive box test. However, while this is certainly the adult interpretation of both intentional and intensional questions, this may not necessarily be the case for children. We believe it is possible to succeed on intentional Q2 but not intensional Q3 simply by appreciating that Heinz only has partial access to information. He sees the ball in the box and does not hear that it is a present but the relation between these two facts (i.e. that they both share the same object of reference) is not represented: in this respect we suggest the child does not represent Heinz's knowledge as partial.

Without the support of such a representation, children cannot make an opaque interpretation of the intensional context which, as we saw above, arises precisely because the way we describe the object of Heinz's knowledge is constrained by what Heinz knows. To labour the point, if the connection between the facts that Heinz knows about the ball but not about the present is not represented, the child's own knowledge that they co-refer will force transparent interpretations of intensional contexts: either term can describe Heinz referring to the object since for the child they do co-refer. This point is not pedantic but crucial, because transparent readings have very different effects upon the meanings of questions 2 and 3.

Intensional Q3 has an opaque reading corresponding to something like: 'does Heinz know that 'there's a present in the box'?', which has the correct answer 'no'. In contrast, a transparent reading would be: 'does Heinz know there's something in the box?' which has the correct answer 'yes', since Heinz has seen inside the box. This is clearly different from intentional Q2 which has the opaque reading: 'does Heinz know that 'the ball is a present'?' (correct answer 'no'), and the transparent reading: 'does Heinz know that thing is a present?' which still has the correct answer 'no'. Thus, whilst for an adult, the natural reading of intentional Q2 contrasts explicitly what we know with what Heinz knows, the answer to the question remains the same under the other possible reading.

This analysis shows the crucial difference between the intentional and intensional questions very clearly, and allows a basic failure to represent the partial nature of mental representations to remain in the running as a reason for children's difficulties. We thus have two candidate explanations. The first suggests that children lack an understanding of the rules governing the substitution of co-referential terms in

intensional contexts, which according to Russell (1992) is a form of specifically-linguistic knowledge. The second locates their problems at the much more fundamental level of the representation of referential relations between reality and the mind. Experiment 2 was an attempt to separate these two explanations.

3. Experiment 2

If children's problems in Experiment 1 were due to a lack of knowledge about the pragmatic substitution rules in talk about mental states, the obvious prediction is that where there is no linguistic intensional context, reasoning about partial knowledge will be no more difficult than any other theory-of-mind task. In contrast, if children have a basic problem representing this consequence of the representational mind, predicting actions arising from partial knowledge will be as difficult as describing it. In Experiment 2 children were asked to predict Heinz's search behaviour on the basis of partial knowledge. This method had two additional advantages over Experiment 1. First, it was possible to make a very exact false-belief style control so that there was no uncertainty about task-specific difficulties. Second, whilst in the language-based situations so far discussed, children's answers to the intensional questions could only be conventionally wrong, in Experiment 2 children's mistakes would be clear errors – not a matter of interpretation but simply wrong.

3.1. Method

The principal test conditions used objects that could be referred to by two definite descriptions (X and Y), only one of which was visually obvious (X). These were paired with objects whose appearance demanded the non-obvious definite description Y. So for example, the dice/rubber from Experiment 1 was paired with a rubber of normal appearance. After the child had examined the objects and the definite descriptions had been agreed, s/he watched as the puppet was allowed visual access only. As in Experiment 1, we checked that children had followed the scenario with intentional Q2: 'does Heinz know that the X is a Y?' Presentation of this question was counter-balanced with pseudo-intensional Q1 'where will Heinz go to look for a Y?' This question was labelled pseudo-intensional, in recognition of the fact that, whilst it was not intensional in itself, a correct answer required the child to represent Heinz's knowledge. This meta-representation would necessarily entail an intensional context. Importantly, each child received both questions within trials. Recall that in the discussion of Experiment 1 we mentioned the possibility that children's extensional errors were merely performance problems and that these might be eliminated if children were prompted to consider the partial representational content. The above design means that half of the children receive just such a prompt, in the form of intentional Q2, before they are asked to predict what Heinz will do.

Earlier, we noted that the handling of ambiguity and intensional contexts both seem to require that the partial nature of representations be represented. We were therefore interested in whether performance on a test that required the child to

recognise ambiguity in an utterance would be related to that on the pseudo-intensional tasks.

3.1.1. Participants

We tested 39 children (19 boys and 20 girls) aged between 4 years 3 months and 5 years 2 months (mean age 4 years 9 months) who were in reception classes, 54 children (27 boys and 27 girls) aged between 5 years 3 months and 6 years 2 months (mean age 5 years 9 months) from year 1 classes, and 37 children (18 boys and 19 girls) aged between 6 years 3 months and 7 years 2 months (mean age 6 years 9 months) from year 2 classes. We shall refer to these as the 4-, 5- and 6-year-olds, respectively. All the children attended the same infant/junior school with a lower middle class/upper working class catchment area in Birmingham, UK, and spoke English as their first language.

3.1.2. Materials

There were three different pairs of items used for the two pseudo-intensional conditions and the intentional control condition respectively. In the pseudo-intensional conditions, one item in each pair had a dual function, only one of which was visually apparent. The dice/rubber from Experiment 1 was used again and was paired with a rubber of normal appearance. The rubbers were contained in separate boxes. The second item was a baby's toy that looked like a ball and contained a bell which was not apparent unless the ball was shaken. This was paired with a typical bell. Rather than boxes, covers that could be removed noiselessly were used to hide these items (see Section 3.1.3). The intentional control condition used a pair of identical pencil sharpeners that could be hidden under two up-turned tins. For the ambiguity task we used a Lego chassis with missing rear wheels that could be fitted with identically-coloured (both red and black) wheels of different sizes. The deceptive box control task was as described in Experiment 1. The puppet, Heinz, was used once more as the co-protagonist in each trial.

3.1.3. Procedure

Children were tested individually with all five items. In the two pseudo-intensional tasks, children always received the item with the dual identity first. With the dice, they were allowed to see the item in its box and the visually-obvious label was agreed before they felt it. If children did not identify it as a rubber spontaneously upon feeling, it was demonstrated. Once agreed, the two descriptions were then reiterated: '...so it looks like a dice and it feels like a rubber...' Next, they looked in the second box and were asked to identify the other rubber. Children had no difficulty with these naming tests. Next, children observed as Heinz looked inside the boxes, now placed apart on the table in front of the child (both the relative spatial position of the two boxes and the order in which the puppet looked was random). As before, Heinz's lack of other perceptual access was emphasised by saying 'now Heinz is looking inside the boxes, but he's not feeling'. With the lids back on the boxes, children were shown a cartoon drawn in pencil and told that Heinz wants to change the mouth on the figure from a frown to a smile, so he needs to find a rubber.

Children were asked either the pseudo-intensional question first:

Q1 ‘where will Heinz go to find a rubber?’

and regardless of their answer, this was followed by the intentional question

Q2 ‘does Heinz know that the dice is a rubber?’,

or they received these questions in the opposite order. This ordering was counterbalanced between child but was consistent within child. Children’s responses invariably took the form of a pointing gesture.

Trials with the ball/bell followed the same general form. This time, children heard and identified the bell before they saw and identified it as a ball. Again, these labels were reiterated with emphasis upon what information was available from seeing versus hearing: ‘it looks like a ball and it sounds like a bell’. Both items were placed on the table in front of the child and individually covered. After Heinz had viewed the items under the covers the child was told that he wanted to make some noise to wake up his friend who was asleep under the table. The pseudo-intensional question then became:

Q1 ‘where will Heinz go to make some noise?’

whilst the intentional question remained

Q2 ‘does Heinz know that the ball is a bell?’

The intentional control task had a similar surface form but was crucially different, because this time Heinz did not know about one of the objects. Children were introduced to one pencil sharpener in the presence of Heinz. They were then told that he wants to hide it under one of the (two up-turned) boxes on the table. Once he had seen the sharpener hidden under one box, Heinz was removed from the table and put out of the way in a bag. Children were then shown a second identical sharpener and it was suggested that they hide it under the other box on the table. It was emphasised that Heinz could not see what we were doing. Heinz then returned to the table, needing a pencil sharpener. Children were either asked the action question:

Q1 ‘where will Heinz go and look for a pencil sharpener?’

followed by intentional knowledge question:

Q2 ‘Does Heinz know that there’s a pencil sharpener in here?’ (pointing to the second box)

or they received the Q2 first and Q1 second.

The ambiguity task was designed to be a relatively easy test of children’s understanding of ambiguity in referential communication, in that children were asked to judge the listener’s knowledge rather than to identify the source of the problem (Robinson and Whittaker, 1987; Sodian, 1988). Children were shown a Lego chassis and told that Heinz had been making this toy, but that ‘...he can’t decide which wheels to put on it’. Children were shown that both sets of wheels fitted, producing a car-like or tractor-like model, then the wheels were hidden under two boxes on the table. Children were told ‘now we’re going to get Heinz to make up his mind, we’re going to ask him to tell us exactly which ones he wants’. Heinz then ‘whispered’ to the experimenter, who relayed the ambiguous message to the child: ‘he says he wants the red and black wheels’. The test question followed:

‘do we know which ones he wants?’

The procedure for the deceptive box task was as described in Experiment 1.

The five trials and their variants were partially counter-balanced. The three pseudo-intensional and intentional trials were always blocked together, but were ordered in the six possible ways. Half of the time, these tasks were preceded by the deceptive box and ambiguity tasks, in that order. On the other occasions, the ambiguity and deceptive box tasks came after the three other tasks, in that order.

3.2. Results and Discussion

There were no effects of task order so all of the orders were combined for further analysis.

Performance on the intentional control task Q1 was not significantly above the chance baseline in the 4-year-olds: 24/39, 61%. Performance on Q2 ('does Heinz know there's a sharpener in here?') was similar with children achieving 26/39 (67%) correct. As expected, the 5-year-olds performed better on Q1: 47/54, 87%, significantly above chance, $P < 0.0001$ by sign test). This was significantly better than the 4-year-olds (χ^2_(1, $n = 93$) = 16.13, $P < 0.001$). The 5-year olds also improved on Q2: 52/54 (96%) (χ2 (1, $n = 93$) = 14.69, $P < 0.001$). There was no significant improvement on this near-ceiling performance between the 5- and 6-year-olds with children in year 2 scoring 31/37 (84%) on Q1 and 37/37 on Q2. In the -5 and 6-year-old groups, five and six children, respectively, were correct on Q2 and incorrect on Q1, while no children showed the opposite pattern. This difference was statistically significant only for the oldest group, $P = 0.032$ by sign test. However, answering Q2 first had no effect upon children's answers to Q1, nor vice-versa: all $P > 0.47$ by chi-square test.

On the deceptive box task 19/39 (49%) of the 4-year-olds were successful (from a 0% baseline), suggesting that the intentional control task was, if anything, a conservative measure of children's ability to distinguish intentional and extensional contexts. As would be expected, 5- and 6-year olds performed near ceiling (47/54 and 34/37 respectively).

Responses to the (intentional) knowledge question (Q2) of the pseudo-intensional task showed the same pattern of change as the intentional control. The 4-, 5- and 6-year olds scored an average of 17/39 (44%), 47/54 (87%) and 36/37 (97%), respectively, across the two tests. Children were classified as having got none, one or both of the knowledge questions correct, and on this basis, children showed significant improvement with age, χ^2 (4, $n = 130$) = 47.5, $P < 0.0001$.

In marked contrast, scores on the action questions (Q1: 'where will Heinz go to...') of each of the two pseudo-intensional tasks were not significantly different from chance in any of the three age groups. For the dice/rubber and ball/bell, scores were: 4-year olds 15/39 (38%) and 15/39 (38%), 5-year olds 27/54 (50%) and 22/54 (41%), 6-year olds 24/37 (65%) and 20/37 (54%) respectively. The fact that the rubber was referred to directly and the bell was only 'to make some noise' made no significant difference to children's performance on the two tests. We therefore considered the two tasks together, classifying children according to whether they got none, one or both of the pseudo-intensional questions correct. The combined performance on the two pseudo-intensional tasks was compared with its expected

distribution on the basis of chance performance. On this more sensitive measure, the 6-year-olds, but neither of the younger age groups, performed significantly above chance (goodness-of-fit χ^2 (2, $n = 37$) = 10.4, $P < 0.01$, other $P > 0.05$). Similarly, when combined performance was considered, there was a significant change in children's performance across the three years (χ^2 (4, $n = 130$) = 12.99, $P = 0.011$).

For both 5- and 6-year olds, the action question (Q1) was significantly harder than both the knowledge question (Q2) and the matched intentional control task (all $P < 0.001$ by sign test). This replicates the disparity of performance between intentional and intensional questions we observed in Experiment 1.

Table 3 details children's combined performance on the pseudo-intensional task according to whether they had the pseudo-intensional or the intentional knowledge question first. The tendency was towards better performance when the intentional question came first, but no order effects approached significance (all $\chi^2 < 1.41$, all $P > 0.494$): asking children first whether Heinz knew the dual identity of the test item did not assist them in predicting where he would search. This suggests that children's difficulty with the action question was not simply due to their failure to recall Heinz's knowledge state. This is indeed striking, because even if one does not normally subscribe to the performance error accounts outlined in the introduction, one might reasonably expect such a relevant prompt to assist children.

In the ambiguity task there was a pattern of improving performance from the 4-year-olds (6/39, 15%) to 5-year-olds (24/54, 44%) to 6-year-olds (23/37, 62%) which was significant (χ^2 (2, $n = 130$) = 17.72, $P < 0.0001$) and in line with other results in the literature (see Robinson, 1994). As we expected, the ambiguity task was consistently harder than the deceptive box task (all $P < 0.001$ by sign test). To compare performance on the ambiguity task with that on the pseudo-intensional tasks, scores for the two pseudo-intensional tasks were combined. There was no

Table 3
Number (percentage of sub-group) of children in the three age groups who got zero, one or both pseudo-intensional questions correct, divided according to whether they received the questions before or after the intentional knowledge question

Order of questions in the pseudo-intensional condition	Number of pseudo-intensional questions correct	Reception (4 years) $n = 39$	Year 1 (5 years) $n = 54$	Year 2 (6 years) $n = 37$
Intentional question first	0	4 (29%)	7 (29%)	3 (19%)
	1	8 (57%)	11 (46%)	4 (25%)
	2	2 (14%)	6 (25%)	9 (56%)
	Mean	6 (43%)	11.5 (48%)	11 (69%)
Pseudo-intensional question first	0	10 (40%)	9 (30%)	7 (30%)
	1	12 (48%)	16 (53%)	6 (29%)
	2	3 (12%)	5 (17%)	8 (38%)
	Mean	9 (36%)	13 (43%)	11 (52%)

significant contingency between the two types of task for the 4-year-olds. For the 5-year-olds there was just significant contingency (χ^2 (2, n = 54) = 6.01, $P < 0.05$). For the 6-year-olds the contingency was non-significant (χ^2 (2, n = 37) = 5.66, $P = 0.059$) largely as a result of noise from the pseudo-intensional task's 50% base line. Following the method in Everitt (1992, pp. 41–44) we were able to reduce this problem by partitioning the χ^2 statistic. First, a contingency test was performed comparing success on the ambiguity task for children who got either none or one of the pseudo-intensional action questions (Q1) correct (χ^2 (1, n = 20) = 0.213). These two categories (that theoretically accounted for 75% of chance responses) were then combined, and compared with the ambiguity results of children who had got both pseudo-intensional action question correct. Contingency between the ambiguity and pseudo-intensional tasks was then significant (χ^2 (1, n = 37) = 5.45, $P < 0.02$), suggesting, together with the results from the 5-year-olds, that the ambiguity and pseudo-intensional tasks may place common demands on children.

To summarise, the results replicate the well-known finding that by 5-years of age, children have little difficulty making the intentional/extensional distinction (as required for success on deceptive box and unexpected transfer tests of false belief, and the intentional control condition in the current study). Performance on the ambiguity task is also consistent with the literature. Set against these benchmarks, the children's performance on the pseudo-intensional task is quite striking (see summary in Table 4). In line with their performance on the deceptive box test and intentional control, 5- and 6-year-olds found the intentional knowledge question (Q2) easy: 'of course Heinz doesn't know that the dice is a rubber; he hasn't felt it!' In contrast, these children were at chance when they had to apply this knowledge and predict where Heinz would look. However, there were significantly more 6-year-old s passing both tasks than expected by chance, which suggests that at least some of these oldest children were fully competent. Additionally, for both the 5- and 6-year-olds there was significant contingency between performance on the pseudo-intensional tasks and the ambiguity task. In the discussion we will argue that this is no

Table 4
Comparison of class performance across five of the experimental conditions

Age group	Intentional control condition, Q1 50% baseline	Intentional Q2 from 'pseudo-intensional conditions (mean)	'Pseudo-intensional' conditions (mean) 50% baseline	Ambiguity condition	Deceptive box condition
Reception (4 years) n = 39	24 (61%)	17 (44%)	15 (38%)	6 (15%)	19 (49%)
Year 1 (5 years) n = 54	47 (87%)	47 (87%)	24.5 (45%)	24 (44%)	47 (87%)
Year 2 (6 years) n = 37	31 (84%)	36 (97%)	22 (59%)	23 (62%)	34 (92%)

Both intentional control and 'pseudo-intensional' conditions have theoretical 50% baselines while empirical evidence tells us that the ambiguity and deceptive box tasks have zero baselines.

coincidence but in fact results from a basic similarity in the representational demands of the two conditions.

4. General discussion and conclusions

In Experiment 1 we found a strong dissociation between children's handling of intensional and intentional contexts and we considered three candidate explanations: a lack of linguistic knowledge about substitution rules in intensional contexts, a performance error that masked underlying competence, and a basic problem representing referential relations. In Experiment 2 we found three lines of evidence to support the last explanation.

First, linguistic knowledge of the kind discussed was irrelevant in Experiment 2, since there simply was no linguistically-grounded intensional context and thus no opportunity to make a non-conventional, transparent interpretation. But this does not eliminate referential problems because, as Searle (1983) stresses, intensionality results from representations only partially capturing their referents, it is independent of the mode in which meta-representation occurs and thus is not a linguistic phenomenon per se. According to Searle, at least some of the problems with intensional contexts that occur in the linguistic realm, are mirrored precisely in mental representation. Therefore, understanding reference between multiple situations must require partitioning of the aspects under which an object is respectively represented. This is exactly what we believe children are failing to do. The only substantial difference between the two main conditions in Experiment 2 is that in the pseudo-intensional task, Heinz's knowledge is partial whilst in the control – as in all of the standard theory-of-mind test formulations – ignorance is total.

The second line of evidence against explanation in terms of linguistic knowledge concerns the observed correlation between the pseudo-intensional condition and the ambiguity task: Ambiguity has nothing to do with the pragmatics of when an opaque or transparent interpretation is appropriate. In contrast, an inability to handle the partial nature of representations (mental or linguistic) could explain the tendency to judge that an utterance adequately discriminates two similar objects when in fact it refers equally to both. The requirement to hold in mind the relation between the fully specified 'real-world' object and its partially-specified representation is common to the handling of both intensional contexts and ambiguity.

What about the possibility of performance problems mentioned in the introduction? Could children's reality bias (e.g. Mitchell, 1996), executive limitations (e.g. Russell, 1996) or economy of effort in selecting explanatory hypotheses (e.g. Fodor, 1992) explain our results? The counter-intuitive finding that answering a question about what Heinz does not know had no significant effect upon children's ability to predict how he will act points away from such explanations. Further, the very close similarity between the intentional control condition (in which children successfully predicted action on the basis of straightforward knowledge and ignorance) and the pseudo-intensional condition, surely makes the suggestion that children fail to select a 'belief-based' hypothesis in the latter difficult to sustain. The necessary refine-

ment, that children who can select a 'belief based' hypothesis fail to select a 'partial belief-based' one when necessary seems somewhat ad hoc.

The results of both Experiments 1 and 2 also appear incompatible with the view that the theory-of-mind tests passed at around 4 years diagnose the development of a theory of meta-representation (e.g. Perner, 1991; Gopnik and Wellman, 1992). Such a theory should be able to cope with intensional contexts as easily as false beliefs. By showing that children's difficulties persist even in very closely-controlled comparisons with other mentalistic tasks, the results presented here suggest that this is not the case.

The developmental account of Chandler (e.g. Carpendale and Chandler, 1996) which withholds a mature understanding of the mind until 6–7, initially seems better equipped to explain the observed pattern of late change. However, we have already suggested that it may fall down on the details. In Section 1, we argued that the difference between a copy theory and an interpretive theory of the mind could not easily explain the dissociation between intentional and intensional questions in Russell's experiment (Russell, 1987). This argument applies equally to the results of our two experiments. Furthermore, in Experiment 2 the close similarity between the intentional control and the pseudo-intensional tasks makes a change in a child's gross theory about how the mind works seem an unnecessarily complex means of explaining the difference in task difficulty. In sum, Chandler's framework has the necessary form but appears to lack the explanatory substance. A sharper instrument than either Perner or Chandler offer seems to be required to explain the dissociation between intentional and (pseudo-) intensional conditions.

How could children's developments be better characterised? In what follows, we speculate how the child's developing cognitive system could give rise to the observed pattern of results.

To handle the referential problems posed by folk psychology requires some kind of representational partitioning: one's representations of mental contents must be separated from one's representations of current, extensional reality. This can be achieved in various different ways, so we summarise first of all the criteria by which we may assess a system for capturing our results. First, a single solution will not do. Amongst other things, an account must accommodate the fact that the ability to handle out-dated representations (such as false beliefs), does not entail the ability to handle their partial nature (as in intensional contexts). Second, we should take account of the fact that the problems of mental representations are but a particular instance of a much more general set of problems, posed for example, by representing any out-dated, hypothetical or spatially-displaced situations (Carnap, 1947; Dowty et al., 1981; Fauconnier, 1985; Perner, 1991). Perhaps a devoted mentalising device is not what we are looking for. Finally, the system should be plausible, not just at a static end point, but developmentally.

Leslie (1987) argues for a syntactic 'decoupling' mechanism, analogous to quotation marks in natural language, that holds normal, reality-oriented representations outside of their usual causal relations for mentalistic calculations. This allows tokens in a language of thought to represent mental contents which may contradict 'reality', without undermining the internal consistency of the system. Leslie claims that this

mechanism underpins children's ability to handle mentalistic phenomena and cites early pretence as evidence that it comes on line at around 18 months of age. Issues of timing aside, Leslie's account of a general-purpose mechanism falls down on our first two criteria. Its presence or absence cannot explain the dissociation found in our results because it deals with the referential problems due to out-dated, hypothetical and partial representations in the same way. It is also devoted to mentalistic calculations and so has nothing to do with the related non-mentalistic referential problems. While it is possible to specify a number of innate syntactic partitioning mechanisms coming on-line to solve different referential problems at various times, this would reduce the parsimony of Leslie's explanation. It might be more appealing instead to exploit the richness of the theoretical and empirical phenomena in a developmental account relying less heavily on 'blind' syntactic mechanisms.

Representing meaning with semantic partitioning may provide the basis for such an account. We take the scheme of Fauconnier as our example though the actual details of implementation are less important for our purposes than the semantic nature of the representations[1] (Fauconnier, 1985, 1997). According to Fauconnier, aspects of meaning are constructed in the mind by the partitioning of content into different 'mental spaces' whose properties lend it different meanings. The same word tokens can create different meanings according to the space to which they are assigned. For example, the possible opaque and transparent readings of propositional attitude reports such as 'Heinz knows there's a ball in the box' correspond to whether the content 'ball' is assigned to a mental space for reality in which reference is extensional or a space corresponding to Heinz's belief in which (for adults at least) co-referential terms are not freely substitutable. Non-mentalistic problems are solved in the same formal way: The alternative readings of 'in 40 years time, my middle-aged daughter will marry' correspond to whether 'middle-aged daughter' is assigned to the mental space for present time or the space for 40 years hence[2].

In contrast to the syntactic approach, in which partitioning is essentially arbitrary and mechanical, the properties that differentiate semantically-partitioned spaces could be called knowledge in a fairly conventional sense of the word. For example, a space for 40 years hence would allow its contents to have entailments that did not change across time, like a person's name, but would suspend those that might change, such as age or hair colour. Developmentally, we could expect this internal structuring to change with experience. Furthermore, it is possible that the same knowledge could structure different spaces with analogous problems, such as representations of past reality and out-dated beliefs.

Consider for example the standard unexpected transfer task (Wimmer and Perner, 1983). To answer correctly that Maxi will look in the old location, the child has to treat the mind as representing an out-dated state of affairs and juggle referential relationships between two temporal contexts. As Perner has noted (Perner, 1991; pp. 54–55) search after invisible displacement requires a similar ability to separate past

[1]Dinsmore, 1991 discusses how his theory of partitioned representations only concerns the highest level of mental representation, leaving open the detailed choices between propositional, analogical and procedural representations.

[2]This example is from Fauconnier (Fauconnier, 1997; Fauconnier, 1985).

from present yet this is achieved around the age of 18 months. As suggested above, this early ability to represent reference across time could provide a basis for representing mental reference across time, so simplifying the task of learning.

The other standard test of false belief, the deceptive box task (e.g. Gopnik and Astington, 1988) can be analysed in a similar way. Here, the child discovers that a sweet box contains, not sweets, but pencils, and then has to predict what an ignorant person will think is inside. To answer correctly, it is necessary to take into account the protagonist's incorrect expectation, in the face of one's own knowledge of reality. Now, the child has to model referential relations between a real and a 'possible' state of affairs, and again there may be a precedent, this time in early pretence. Hence, the ability of 4-year-olds to treat people's minds as representing out-of-date or hypothetical states of affairs may be a development of their earlier abilities to represent times, possibilities and places other than the here and now.

Although both false-belief tasks involve analysing a mental representation, neither requires the child to represent the fact that mental representations are only partial. Children can succeed on these tasks even if the partitioned space corresponding to 'Heinz's belief' does not restrict the characterisation of its contents to the aspects of which Heinz is aware. Failure in the intensional and pseudo-intensional tasks implies the absence of this ability to hold a representational content under particular descriptions. This time there is no obvious non-mentalistic precedent in children's earlier abilities, so it is perhaps not surprising that this is harder.

Considering children's developing ability to handle mental representation in this way, in terms of changes in the processes of semantically-based partitioning, poses in a novel way the question of the role of experiences with language and social behaviour in the development of adult folk psychology. This approach emphasises representational competence along with the acquisition of knowledge. The obvious questions arising from the current study concern the precise nature of children's understanding when they appear to follow Heinz's partial information access but fail to represent adequately the resulting partial knowledge. How does this provide a basis for children to move forward towards adult-like performance and what processes does this entail?

Acknowledgements

The research reported was carried out in partial fulfilment of the degree of PhD for the first author, who was supported by a studentship from the Biotechnology and Biological Sciences Research Council, UK

References

Carnap, R., 1947. Meaning and Necessity. University of Chicago Press, Chicago.

Carpendale, J., Chandler, M., 1996. On the distinction between false belief understanding and subscribing to an interpretive theory of the mind. Child Development 67, 1686–1706.

Chandler, M., 1988. Doubt and developing theories of mind. In: Astington, J.W., Harris, P.L., Olson, D.R. (Eds.), Developing Theories of Mind. Cambridge University Press, New York, pp. 387–413.

Chandler, M., Boyes, M., 1982. Social cognitive development. In: Wolman, B. (Ed.), Handbook of Developmental Psychology. Prentice-Hall, Englewood Cliffs, NJ, pp. 387–402.

Chandler, M., Helm, D., 1984. Developmental changes in the contribution of shared experience to social role taking competence. International Journal of Behavioural Development 7, 145–156.

Dinsmore, J., 1991. Partitioned Representations. Kluwer, Dordrecht.

Dowty, D.R., Wall, R.E., Peters, S., 1981. Introduction to Montague Semantics. Reidel, Dordrecht.

Everitt, B.S., 1992. The analysis of contingency tables, second ed. Chapman and Hall, London.

Fauconnier, G., 1985. Mental Spaces: Aspects of Meaning Construction in Natural Language. MIT Press, Cambridge, MA.

Fauconnier, G., 1997. Mappings in Thought and Language. Cambridge University Press, Cambridge, UK.

Flavell, J.H., Speer, J.R., Green, F.L., August, D.L., 1981. The development of comprehension monitoring and knowledge about communication. Monographs of the Society for Research in Child Development. Serial No. 192. University of Chicago Press, Chicago, IL, USA.

Fodor, J.A., 1992. A theory of the child's theory of mind. Cognition 44, 283–296.

Gopnik, A., Astington, J.W., 1988. The development of children's understanding of representational change. In: Astington, J.W., Harris, P.L., Olson, D.R. (Eds.), Developing Theories of Mind, first ed., vol. 1. Cambridge University Press, Cambridge, UK, pp. 193–206.

Gopnik, A., Wellman, H.M., 1992. Why the child's theory of mind really is a theory. Mind and Language 7, 145–171.

Leslie, A.M., 1987. Pretence and representation: the origins of 'theory of mind'. Psychological Review 94, 412–426.

Leslie, A.M., 1994. Pretending and believing: issues in theory of ToM. Cognition 50, 211–238.

Mitchell, P., 1994. Realism and early conception of mind: a synthesis of phylogenetic and ontogenetic issues. In: Lewis, C., Mitchell, P. (Eds.), Children's Early Understanding of Mind: Origins and Development, first ed. Lawrence Erlbaum, Hove, E. Sussex, pp. 355–381.

Mitchell, P., 1996. Acquiring a conception of mind. A review of psychological research and theory. Psychology Press, Hove, E. Sussex.

Mitchell, P., Lacohee, H., 1991. Children's early understanding of false belief. Cognition 39, 107–127.

Mitchell, P., Robinson, E.J., Nye, R.M., Isaacs, J.E., 1997. When speech conflicts with seeing: young children's understanding of informational priority. Journal of Experimental Child Psychology 64, 276–294.

Moore, C., Jarrold, C., Russell, J., Lumb, A., Sapp, F., MacCallum, F., 1995. Conflicting desire and the child's theory of mind. Cognitive Development 10, 467–482.

O'Neill, D.K., Astington, J.W., Flavell, J.H., 1992. Young children's understanding of the role that sensory experiences play in knowledge acquisition. Child Development 63, 474–490.

Perner, J., 1991. Understanding the representational mind, first ed. MIT Press, Cambridge, MA.

Perner, J., Davies, G., 1991. Understanding the mind as an active information processor: do children have a 'copy theory of mind'?. Cognition 39, 51–69.

Pillow, B.H., 1989. Early understanding of perception as a source of knowledge. Journal of Experimental Child Psychology 47, 116–129.

Pillow, B.H., 1993. Pre-school children's understanding of the relationship between the modality of perceptual access and knowledge of perceptual properties. British Journal of Developmental Psychology 11, 371–389.

Pratt, C., Bryant, P., 1990. Young children understand that looking leads to knowing (so long as they are looking into a single barrel). Child Development 61, 973–982.

Robinson, E.J., 1994. What people say, what they think and what is really the case: children's understanding of utterances as sources of knowledge. In: Lewis, C., Mitchell, P. (Eds.), Children's Early Understanding of Mind: Origins and Development, first ed. Lawrence Erlbaum, Hove, E. Sussex, pp. 355–381.

Robinson, E.J., Mitchell, P., 1992. Children's interpretations of messages from a speaker with a false belief. Child Development 63, 639–652.

Robinson, E.J., Mitchell, P., 1994. Children's judgments of ignorance on the basis of absence of experience. British Journal of Developmental Psychology 12, 113–129.
Robinson, E.J., Mitchell, P., 1995. Masking of children's early understanding of the representational mind: backwards explanation versus prediction. Child Development 66, 1022–1039.
Robinson, E.J., Robinson, W.P., 1982. Knowing when you don't know enough: children's judgments about ambiguous information. Cognition 12, 267–280.
Robinson, E.J., Thomas, G.V., Parton, A., Nye, R., 1997. Children's overestimation of the knowledge to be gained from seeing. British Journal of Developmental Psychology 15, 257–273.
Robinson, E.J., Whittaker, S.J., 1987. Children's conceptions of relations between messages, meanings and reality. British Journal of Development Psychology 5, 81–90.
Russell, J., 1987. 'Can we say...?' Children's understanding of intensionality. Cognition 25, 289–308.
Russell, J., 1992. The theory theory: so good they named it twice?. Cognitive Development 7, 485–519.
Russell, J., 1996. Agency: Its Role in Mental Development. Hove, E. Sussex, UK.
Russell, J., Jarrold, C., Potel, D., 1994. What makes strategic deception difficult – the deception or the strategy?. British Journal of Developmental Psychology 12, 301–314.
Searle, J., 1983. Intentionality, first ed., vol. 1. Cambridge University Press, Cambridge, UK.
Sodian, B., 1988. Children's attributions of knowledge to a listener in a referential communication task. Child Development 59, 703–718.
Taylor, M., 1988. Conceptual perspective taking: children's ability to distinguish what they know from what they see. Child Development 59, 703–718.
Wellman, H.M., 1990. The child's theory of mind. MIT Press, Cambridge, MA.
Wimmer, H., Perner, J., 1983. Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. Cognition 13, 103–128.