Young children's understanding of other people's inferences

ANGELA VAROUXAKI[1]
NORMAN H. FREEMAN
University of Bristol, United Kingdom

ABSTRACT When do children begin to understand inference as a source of knowledge? Estimates in the literature have varied between 4 and 6 years of age. We set out to determine whether previous research has underestimated children's performance. We used mathematical syllogisms to make it salient to the children that a sum needed to be worked out. One hundred one children were tested and their data were analysed using chi-square. The results indicated that about two thirds of the 5- and half of the 4-year-olds had the ability to attribute inferential knowledge to another person. In contrast with previous research our children were able to explicitly justify their correct answers. Finally, no evidence was found to support the idea that children substitute a 'not seeing is not knowing' default rule for inference.

Key words: Inference, theory of mind, young children.

Introduction

Over the past twenty years there has been an international fascination with understanding children's knowledge about people's mind: At around the age of four children benefit from a conceptual shift in which they come to know that reality can be misrepresented, as well as be represented accurately, in the mind (Perner, 1991; Wellman, 1988; Wellman 1990). For example, at this age children understand that people can give different interpretations to one and the same event. Such knowledge is referred to as a core accomplishment in a representational theory of mind (Perner, 1991; Wellman, 1990). According to the theory view the contents of the mind are tools for predicting and explaining other people's as well as one's own behaviour. In other words, the key to interpreting what people do is to understand their mental lives: how people think, what they desire, what their ambitions and intentions are, and so on. Once children start to understand the mind, armed with a primitive theory, they become committed to a life-long career as common-sense psychologists.

An important distinction needs to be made between the first and a second phase of the development of children's theory of mind. The first phase consists in a grasp of perception and communication as the causal origins of knowledge, thoughts and beliefs (Flavell, 1983; Perner, 1991). For example, someone who sees a cake inside a box thereby *knows* what is in the box, and if that person then fails to observe her friend stealing the cake she will be left with a *false belief* that the cake is still in the box. Children can grasp

1. We are grateful to the children and their teachers for their willing participation in this study. Many thanks to D. C. Cooke for assisting with the collection of data. We also thank K. Maridaki-Kassotaki for her useful comments on an earlier version of this manuscript.

Address: Angela Varouxaki, Department of Psychology, University of Bristol, 8 Woodland Road, Bristol, BS8 1TN., U.K. E-mail: a.varouxakis@bristol.ac.uk

such insights if they understand the causes and sources of knowledge, seeing it as contingent upon having *informational* access to the relevant facts (Wimmer, Hogrefe, & Perner, 1988).

Yet, there is a limit to how far such an understanding can get a child in working out what people know. The fact is that people come to know things not just by direct observation and communication. Knowledge can also be acquired by manipulating the relevant evidence to think through problems to a genuinely new piece of information. For example, that is what every police detective does in thinking through where the clues of a crime lead. In brief, not just observation and communication but *inference* is a source of knowledge. The suggestion is that the second phase of development is to grasp that people have minds which actively generate knowledge by means of making inferences. The mind is a powerful inferential engine.

There is a vast step between using inferences oneself and understanding that other people make inferences as a way of gaining knowledge. Halliday (1977) showed that children as young as three years of age were able to use inference for themselves as a means of achieving a goal. In his experiment children were trained on two behavioural sequences - putting (a) a square and (b) a triangle through the correct holes in an apparatus; and were then tested in a situation in which these two behaviour sequences had to be put together appropriately in order to obtain a sweet. Pears and Bryant (1990) found that four-year-old children could make transitive inferences (A>B, B>C, C>D, D>E, therefore B>D) about spatial position without any prior training. What the above studies do not tell us, however, is whether children are aware of where and how they got their knowledge of how to get a sweet or how B relates to D. Children understand the role of inference in the origin of knowledge when they grasp the causal connection between (a) access to the premise information and (b) the resulting epistemic state. To adopt Wimmer, Hogrefe and Sodian (1988), correct assessment of a child's own knowledge indicates the mere functioning of inference as a source of knowledge, whereas correct assessment of another person's inferential knowledge indicates an understanding of the fact that inference is a source of knowledge.

The first study that directly addressed the issue of children's understanding of inference as a source of knowledge ran as follows. Sodian and Wimmer (1987) allowed the child and the other (actually a doll) to look into a container with balls all of one colour (red) or two colours (red and blue). A red ball was then moved from the container into an opaque bag in full view of *either* the child *or* the doll. Access to the transfer of the ball was manipulated by closing a screen, leaving the container on either the child's or the doll's side of the screen. Each child participated in four tasks created by crossing one or two different colour balls with visible or invisible transfer. The critical questions addressed to the children were: Do you know what the colour of the ball in the bag is? and Does [name of the doll] know what the colour of the ball in the bag is?. In the first of a series of experiments, 14 out of 64 children overestimated their own knowledge by incorrectly claiming that they knew the colour of the ball in the opaque bag despite not actually seeing the ball being taken out of the two-colour container. Of the children who were incorrect for the other (doll), there were 87% of errors in judging the other's knowledge which were neglect of the other person's use of inference. That is, most of the errors were to ignore the fact that when the other knew that there were only red balls in the container, the child could infer that the ball taken out and put into the bag must be red, without having actually to see the transfer being made. Sodian and Wimmer (1987) suggested a conceptual deficit to account for their data. But the procedure was not entirely satisfactory. It is so easy for the child automatically to work out whether she knows that a ball taken from a set of red balls will still be red when it is moved to a bag, that she might not realise that any inference is involved at all in the situation.

Keenan, Ruffman and Olson (1994) suggested that some of the variance that explains children's difficulty with inference is due simply to forgetting

the initial-state information. Keenan et al. (1994) introduced an *inference salient* condition in two ways: first, children were shown a picture of the initial-state after the transfer of the ball to the opaque bag. Second, they were reminded that the other person knows the colour of the balls too. Note that *either* the child *or* the other person was given enough information to solve the inference, but *not* both. In the condition where the other person was deprived of access to the end-state information success was high in judging that the other did not know what was in the bag: 69%, 75%, 88% and 100% at ages 3, 4, 5, and 6 years respectively. Of those children, neglect of inference as a source of knowledge in the converse condition was shown by 69%, 47%, 36% and 20% respectively. Keenan, Ruffman and Olson (1994) concluded that a fledging understanding of inference as a source of knowledge begins to develop sometime around 4 years of age and takes a long time to become consolidated.

The present study

The aims of the present study were to determine (a) the extent to which 4- to 5-year-old children understand inference as a source of knowledge using simple arithmetic and (b) the level of explicitness of 4- to 5-year-old children's understanding of inference.

It is possible that pre-schoolers acquire a valuable insight that direct observation causes knowledge and might assume that the obverse holds so that not seeing equals not knowing (Perner, 1991): that would lead them to overlook cases where someone else can infer a state of affairs without direct looking. Keenan, Ruffman and Olson's (1994) salience condition focused on getting the child to recall the initial-state information, rather than on the underlying inferential procedure that the children need to employ to solve the task. What we need is a design where it is salient to the children that the task is inferential: simple arithmetic might draw to children's attention that something needs to be worked out. This is because simple arithmetic, as opposed to the experimental procedures used by Sodian and Wimmer (1987) and Keenan, Ruffman and Olson (1994) has long been used in an educational context and children are familiar with it. Do children who have a cue that a computation is being worked out credit another mind with the output? Our solution was to have the child and another person count a number of things, then to hide some of the things, allow the child and the other to recount the reduced number left out, and find out whether the child realises if she and the other person know how many of the things had been taken away and hidden.

A second point of interest was to determine how explicitly children were aware of using inference. Karmiloff-Smith (1992) argued that once children achieve practical behavioural mastery of something they normally progress towards a theory of what they are doing, eventually becoming aware of their theory. Children's yes/no answers to the inference question can tell us whether the child credits the other person with using inference as a source of knowledge. But it cannot tell us anything about how deep the children's understanding is. For that, we needed to obtain explanations from the children about how they knew what was in the other person's mind. Once children consciously grasp the reason why someone possesses or lacks inferential knowledge, and can put it into words, they can be credited with a deep understanding of the mind as a generator of knowledge. In Keenan, Ruffman and Olson (1994) study only a negligible minority of children were able to justify their answers. The authors argued that children's ability to justify inference lags behind their ability to understand the phenomenon itself, and the experimenter did not probe further into the matter. So all we can conclude from existing data is that there are some tests which some 4-year-olds pass, and we cannot make a statement about the speed of development of the children's explicit grasp of why they make the judgements they do.

To sum up: there is only little evidence in the literature that bears on the question of when children understand inference as a source of

knowledge; and the task that has been used was not designed to make it clear to children that inferences are involved in working out the answer (Keenan, Ruffman, & Olson, 1994; Sodian & Wimmer, 1987). In the experiment that follows we (a) tested the hypothesis that an arithmetic task would enhance children's performance, because it makes it clear that an inference needs to be worked out. Further, no-one has yet assessed children's explanations, so it is not yet possible to make a statement about the speed of development of their understanding of inference from behavioural mastery to a fully explicit concept. The experiment below rectifies that. We (b) explored 5- and 4-year-old children's ability to justify their assessments of their own and the other person's knowledge. The degree of explicitness of children's answers will bear on the speed of development of their understanding of inference from behavioural mastery to a fully explicit concept.

Method

Participants

Participants were 101 children drawn from two primary state school in a middle class residential area of Avon. There were 44 four-year-olds in the reception class ranging in age between 51 and 59 months (21 girls and 23 boys, M=56 months, and SD=2.6 months); and 57 five-year-olds in year-1 class, with range 61 to 71 months (20 girls and 37 boys, M=66 months, and SD=3 months).

Materials

A metal saucepan and its lid acted as a container into which a number of carrots (taken from a pile of five in each session) were transferred. Two cassette recorders were used to record the conversations (a) between the children and the experimenter and (b) between the children and the experimenter's assistant. An opaque bag containing paper plates was placed under the table. Finally, a pot of salt was hidden at a pre-arranged place in the test room.

Procedure and design

The experimenter and her assistant were introduced to the children by their teacher. Children were tested in pairs in a quiet room close to their classrooms. Prior to testing they engaged in a short conversation with the experimenter to ensure that the children were confident enough to give both yes and no answers to questions.

The children were then asked to help the experimenter pretend to do some cooking. Five carrots were placed on the table, and children were asked to count the carrots aloud, along with the experimenter. Then, one or both of the children were asked to look for some paper plates in the opaque bag under the table. Whilst the participant(s) went under the table some of the carrots (the number varied between 1 and 4 in each session) were moved by the experimenter into the saucepan and the lid was replaced. Hence, either neither or one of the children witnessed the transfer. When appropriate the experimenter made the statement: I want [X] to stay with me and see me putting some carrots into the saucepan. This was done to ensure that the child under the table was fully aware of the other having perceptual access to the number of carrots transferred. *Thus, the first variable introduced in this study was having or not having direct observation of the transfer of the carrots into the saucepan.* Under all conditions, the experimenter verbally communicated to both children that some of the carrots were being transferred into the saucepan. Then, the children were asked to count the number of carrots still left in full view on the table.

There were three levels of the variable of whether or not children had perceptual access to the transfer of the carrots into the saucepan. Level (I-I) was purely inferential since both children were denied perceptual access to the

transfer but could infer the correct number of carrots inside the saucepan. In (P-I), the target child had visual access to the number of carrots transferred into the saucepan whereas the other child was granted only inferential access. Here, the target children were implicitly challenged to refrain from applying a *not seeing is not knowing rule*. Finally, level (I-P), in effect a control, comprised subjects who were blocked from seeing the transfer of the carrots (but could still work the number out via inference) whereas the other child was granted visual access to the contents of the saucepan.

Next, one of the children was asked to help the experimenter's assistant find the pot of salt hidden in the room. *This introduced a second variable, for now, half of the children were still in full view of the array of carrots that remained on the table whilst the other half were denied any memory aid as to the preceding transfer.* Having to retrieve the final array placed extra demands on children's memory. It was expected that having the final array in full view would strengthen children's performance. There were six conditions created by crossing perceptual access to the transfer (either the target, the other, or neither child had perceptual access to the transfer) and perceptual access to the final array of items (either the target or the other child had the final array in full view).

Both participants were then questioned (either by the experimenter or her assistant) about the other child's knowledge: Does [X] know how many carrots are in the saucepan?, and if yes, How many?. Children were also asked to justify their answers, i.e.: How does she know that?/Why doesn't she know that?. Finally, to control for their ability to solve the inference, children were questioned about their own knowledge of the contents of the saucepan: Do you know how many carrots are in the saucepan?; How many?; How do you know that? or Why don't you know that?.

Children who failed the inference question were asked two further control questions: Do you remember how many carrots were on the table when we first counted them? and How many carrots are on the table now?. This safeguarded against memory failures that could had obscured performance. Recall that we were interested in whether or not children can attribute inferential knowledge to another person; the above questions controlled for the possibility that children's misattributions or neglect of another person's inferences stemmed from failure to recall the premise information. Children who failed the memory control questions were re-tested with the above procedure repeated.

Results

Data analysis based on children's yes and no answers

The experimenters were able to establish rapport with all the children, who happily engaged in the task situation. The number of male and female children were roughly balanced in the present study and we have no reason to suspect an effect of sex upon performance. Data were first analysed on the basis of childrens' yes and no answers to the inference questions: Does [X] know how many carrots are in the saucepan? and Do you know how many carrots are in the saucepan? (recall that if the child did not spontaneously come up with the number, she was asked how many? to ensure that she could truly solve the inference). The correct answer was to say yes to the question regarding the other person's and one's own knowledge because both children counted the carrots before and after the transformation. The results from only five 5-year-olds and three 4-year-olds had to be discarded because they failed the question: Do you know how many carrots are in the pot?. Only two 4-year-olds failed the control memory questions and were re- tested. Scrutiny of their data showed nothing unusual and were thus included in the analysis. The distribution of data across conditions is shown in Table 1.

Preliminary analysis of the data showed that whether or not the final array of objects was in full view had no effect on the data. See Table 2.

From Table 2 there was no discernible

Table 1
The frequencies of 5- and 4-year-olds yes/no responses across all experimental conditions

	Five-year-olds		Four-year-olds	
Condition	Yes	No	Yes	No
I-I/ Final array in full view	6	2	2	6
I-I/ Final array to remember	6	2	3	4
P-I/ Final array in full view	3	5	4	4
P-I/ Final array to remember	3	3	4	2
I-P/ Final array in full view	7	2	1	4
I-P/ Final array to remember	10	3	5	2

KEY: I-I: Self Inference - Other Inference
P-I: Self Perception - Other Inference
I-P: Self Inference - Other Perception

Table 2
The frequencies of 5- and 4-year-olds yes/no responses in the conditions where either self or other had perceptual access to the final array of objects

	Five-year-olds		Four-year-olds	
Condition	Yes	No	Yes	No
Final array in full view	16	9	7	14
Final array to remember	19	8	12	8

difference between the childrens who had in full view and those who had to remember the final array of objects in whether or not they correctly assessed the other person's knowledge [x^2 (52)=.23, p=.62 and x^2(41)=2.9, p=.08 for the 5- and 4-year-olds respectively]. Consequently, data were collapsed across this variable for the next analyses.

Table 3 shows the data distributed across three conditions: I-I, P-I and the control I-P. Pooling across all three conditions, the older children did significantly better than the younger [x^2 (93)=4.1, p=.042]. As can be seen in Table 3, 35 out of 42 (67.3%) five-year-olds but only 17 out of 41 (41.5%) four-year-old attributed knowledge to the other person. In fact, while the younger children performed at chance (binomial p>.998), five-year-olds' performance was significantly better than chance (binomial p<.001). So, overall we have a developmental phenomenon: in response to a simple arithmetic task over two-thirds of the 5-year-olds demonstrated an understanding of inference as a source of knowledge. Yet the 4-year-old children did not appear to benefit from the use of arithmetic to make it clear that an inference is involved in working out the answer. However, the issue of chance responding will be returned to shortly when we look at children's explanations.

In the next step we examined 5- and 4-year-olds performance on each of the experimental conditions separately. There were significant differences in performance on the I-I condition

Table 3
The frequencies of 5- and 4-year-old childrens yes/no responses across conditions

	I-I		P-I		I-P		Total	
	Yes	No	Yes	No	Yes	No	Yes	No
Five-year-olds	12	4	6	8	17	5	35	17
Four-year-olds	5	10	8	6	6	6	19	22

KEY: I-I: Self Inference - Other Inference
P-I : Self Perception - Other Inference
I-P : Self Inference - Other Perception

[x^2 (31)=5.4, p=.02]. The data in the control I-P condition was in the same direction but did not quite reach significance (Fishers Exact p=.108). In the P-I condition there were clearly no differences between 5- and 4-year-olds performance [$x^2(28)=.58$, p=.44]. Taken together these data suggest that given a task that is simple enough most 5-year-olds but few 4-year-olds acknowledge another person's inferences. The difference in perspectives in the P-I condition puzzled the children and impeded their ability to attribute inferential knowledge to another person compared with the I-I condition, even when the other was knowledgeable through inference.

In addition to the *quantitative* differences described above, *qualitative* differences were also detected between the 5-year-olds and the younger children on the basis of their yes/no responses. In the I-I and the control P-I conditions, 12 out of 16 (75.0%) and 17 out of 22 (77.2%) five-year-olds respectively correctly attributed knowledge to the other person. Yet, in the P-I condition only 6 out of 16 (42.9%) of the 5-year-olds were correct. The differences in performance between the I-I and P-I; and between the P-I and I-P conditions were significant [$x^2(42)=4.24$, p=.039 and $x^2(42)=.39$, p=.036 respectively]. Five-year-old children appeared to benefit from having themselves to work out the inference prior to assessing the other person's knowledge. In contrast, the 4-year-olds' performance did not vary much across conditions. Chi-square tests showed that the number of 4-year-old children correctly assessing the other person's knowledge was not significantly different between the (I-P) and (I-I); (I-P) and (P-I); nor (P-I) and (I-I) conditions [$x^2(42)=.76$, p=.38; $x^2(42)=1.65$, p=.20; and, $x^2(42)=.13$, p=.71 respectively].

In sum, the results of the present study indicate that under appropriate conditions children come to attribute inferentially gained knowledge to another person sometime between the age of four and five. Such evidence contradicts the pioneering research of Sodian and Wimmer (1987) who concluded that it was not until the age of six that children acquire an understanding of inference as a source of knowledge, and it is consistent with Keenan, Ruffman, and Olson's research. Yet, across conditions, about two thirds of our 5-year-olds and about half of the 4-year-olds attributed inferentially gained knowledge to the other person. What is striking even about the younger children is that they did not systematically apply a *seeing is knowing* rule. If that were the case children in the I-P condition should always conclude that the person who looked inside the saucepan is knowledgeable about its contents.

Analysis of data reconsidered: Children's justifications

Recall that we also asked children to justify their assessment of the other person's state of

knowledge. Both the experimenters and an independent judge classified children's answers to the questions How does [X] know that? or Why doesn't [X] know that?. Since children at the age of 4- to 5-years are limited in terms of their vocabulary and Keenan, Ruffman and Olson (1994) found little evidence of their ability to justify another person's inferential knowledge we set the following minimum criterion: Children had to refer to at least one of the arithmetic premises to be counted as offering an inferential justification. Thus, in the I-I and P-I conditions where the other person knows via the use of inference, examples of justifications that counted as inferential included «There were five and then three left» or «She worked it out because there is only one left on the table». With the exception of the I-P condition perceptual justifications were not deemed appropriate (recall that the I-P was a control condition where the other person could know through perception). These took the form: «She didn't see; the top was on the pan» or «We didn't count how many are in the pan». Other unconvincing justifications either made reference to fantasy or were ambiguous: «Everybody knows that mum, dad, everybody», or «She is stupid». Finally, a number of children failed to justify their answers either by remaining silent or by saying «I don't know».

The inferential justifications provide clear evidence for children's understanding of inference as a source of knowledge. There was no discernible difference between the 35 five-year-olds and the 19 four-year-olds who correctly assessed the other person's knowledge on the basis of their yes/no answers in whether or not they offered an inferential explanation (Table 4). Indeed, 17 out of these 35 (68.6%) five-year-olds and 15 out of the 19 (78.9%) four-year-olds [$x^2(54)=.66$, $p=.41$] had given inferential explanations. There was only one 5-year-old and two 4-year-olds who failed to attribute knowledge to the other person yet went on to offer an inferential explanation. These children can be

Table 4
Frequencies of 5- and 4-year-old children's explanations for their yes/no answers to the inference question across conditions

		Yes/no response across conditions					
		I-I		P-I		I-P	
Age	Explanation	Yes	No	Yes	No	Yes	No
Five-Year-Olds	Inferential	10	0	4	0	10	0
	Perceptual	0	0	0	3	1	0
	Don't know	2	3	2	4	6	5
	Other	0	1	0	1	0	0
Four-Year-Olds	Inferential	3	0	7	0	5	0
	Perceptual	0	2	0	0	0	1
	Don't know	2	7	0	5	1	3
	Other	0	2	1	2	0	0

KEY: I-I: Self Inference - Other Inference
P-I: Self Perception - Other Inference
I-P: Self Inference - Other Perception

considered to have an explicit awareness of the importance of inference as a source of knowledge, even if they did not spot that the other person actually used inference in the situation.

Children's explanations also suggest that children who failed to attribute inferential knowledge to the other person did not substitute a «*not seeing is not knowing rule*». When asked to justify their answers these children either remained silent or merely said «I don't know». Interestingly, only one child used a perceptual explanation in the (I-P) control condition despite the fact that it would have been perfectly appropriate to refer to the other person's perceptual access. Instead, most children used inferential explanations.

In sum, 24 out of the 52 (46.2%) five-year-olds tested and 15 out of the 41 four-year-olds (36.6%) showed respect for another person's inference. The 46.2% and 36.6% correct performance of the 5- and 4-year-olds is no chance responding because not only did these children correctly assess the other person's knowledge on the basis of their yes/no answers but also they clearly articulated the necessity of access to the premise information in order to derive an inference. Children's ability to justify their answers to the questions about the other person's knowledge was somewhat surprising, given findings by Keenan, Ruffman and Olson (1994). In their study only 8 out of the 72 children tested demonstrated an explicit understanding of inference.

Children's justifications provide added support to our claim that children develop an understanding of inference as a source of knowledge sometime around their fifth birthday and suggests that a good many 4-year-olds also have some insight into the causal relation between inference and knowledge.

Discussion

The results of the present investigation cast severe doubt on Sodian and Wimmer's (1987) finding that it is not until the age of six that children begin to understand inference as a source of knowledge.

First, we agree with Keenan, Ruffman and Olson (1994) that an understanding of inference as a source of knowledge begins to develop sometime around the fourth year of a child's life. Some two thirds of our 5-year-olds and half of the 4-year-olds correctly assessed that when another person had access to both the initial and second array of objects s/he was knowledgeable about the number of items transferred. Our use of a paradigm which was modelled on simple arithmetic perhaps made it clear to the children that something needs to be worked out, acting like a catalyst and enabling them to demonstrate an understanding of inference.

Second, our results extend those of Sodian and Wimmer (1987), and Keenan, Ruffman and Olson (1994) by suggesting that not only do some 4- and 5-year-old children have an understanding of the relationship between inference and knowledge but also that such an understanding is highly explicit. Nearly half of our 5-year-olds and nearly forty percent of the younger children showed inference respect using inferential justifications to explain their answers. Sodian and Wimmer did not delve into children's justifications. In Keenan, Ruffman and Olson's (1994) study merely 8/72 children tested were able to justify their answers. This is despite comparable performance between Keenan, Ruffman and Olson's (1994) study and the present study at the level of yes/no responses which suggests that the tasks were of comparable difficulty. Our data also conform to Karmiloff -Smith's (1992) model which suggests that explicit awareness succeeds behavioural mastery of a task. The development of an explicit understanding of inference *rapidly* followed mastery of the task on the basis of children's yes/no responses. What our data do not speak to, however, is when children's understanding of inference reaches ceiling. Testing some older 6- and even 7-year-old children should correct that defect.

A third point to be made from our data is the rarity of appeal to perceptual explanations.

Sodian and Wimmer (1987) suggested that 4- and 5-year-olds were dominated by «a seeing is knowing rule» and so neglected the importance of inference. Overall, only about a sixth of our children who neglected inference as a source of knowledge fell back on a default «not seeing is not knowing rule». Instead, they lacked the ability to explicitly justify their answers altogether. This consistency suggests that the children who neglected inference genuinely lacked an understanding of the causal connection between inferential access and the resulting knowledge. In contrast with Sodian and Wimmer's (1987) suggestion, our children did not indiscriminately draw upon an empiricist theory of knowledge.

One limitation of the above findings is that our children were not assessed for standard false belief competence. This is restrictive in two respects: first, we have no direct means of establishing that our sample of children was a representative one in terms of their theory of mind development. Second, it does not allow us to explore a possible connection between children's performance on the inference task and general theory of mind development. Keenan, Ruffman and Olson (1994) found that children's performance on the inference trials lagged behind that on the false belief task. From a developmental perspective this finding suggests that understanding inference as a source of knowledge is a more advanced acquisition than the development of the ability to metarepresent.

Another limitation of the present study is that the child and the other person were always knowledgeable (either via perception or inference) about the number of items transferred into the container. One might argue that our children were merely imputing their own knowledge of the contents of the saucepan to the other person. It should be noted, however, that the analysis of their justifications should have picked up any children who responded egocentrically. Subsequent research needs to further delve into the above issues.

The question arises of how general is children's understanding of inference as a source of knowledge. Is such understanding limited to the specific inferential task imposed by the experimental design? Are children able to generalise their insight that knowledge is contingent upon inferential access across a number of different situations? Sodian and Wimmer (1994) and Keenan, Ruffman and Olson (1994) used colour and object identity tasks. We used simple mathematical syllogisms. What we need is a design that directly compares children's performance on the two paradigms, controlling for the salience of the inferential procedure necessary to solve the task.

Finally, what determines the onset and consolidation of this phase of the child's theory of mind? A traditional way to study such a question is to ask whether the same factors that influence an earlier phase continue to influence a later phase or whether a qualitative shift to new factors occurs. Research in the United States, Canada, and a number of Western European countries (see Perner, 1991; Wellman, 1990) has focused on some basic aspects of children's representational theory of the mind: children's grasp of the distinction between the mind and the world; the difference between beliefs and desires; and the causal connection between beliefs/desires and action. Yet, such intellectual advances do not occur in a social vacuum. The social context within which children grow up affects the speed of development of a child's theory of mind, depending on factors such as the number of siblings and number of adult kin, as well as the quantity of daily interaction with other adults and older children (Perner, Ruffman, & Leekam, 1994; Lewis, Freeman, Kyriakidou, Maridaki-Kassotaki, & Berridge, 1996). A continuing value accorded to the extended family in many parts of Greece have made it possible to identify a greater range of social influences on the speed of development of the child's theory of mind than can readily be identified in small nuclear families (Lewis et al., 1996). Basically, the more interaction the child has with more knowledgeable people, the faster the child's development in the early phase of a theory of mind. The question now on the agenda is whether that is true of the next phase of the children's development (including an understan-

ding of inference). We suggest that that question can best be answered by adopting the same research strategy as has proved so useful in the first phase. That is, to develop a test of advanced mentalistic competence in nuclear-family children, and then to extend its use to extended families to identify the social factors that control or contribute to children engaging with that second phase of development.

Summary

In sum, in the present paper we set out to determine whether previous research on children's understanding of inference as a source of knowledge underestimated their performance. We found that understanding of inference begins to develop sometime around the age of four. This was consistent with Keenan, Ruffman and Olson (1994) but not with Sodian and Wimmer (1987). Further, the claim that children substitute a not seeing is not knowing rule was not substantiated in the present study. This is not to say that children might not be using such a rule but this alone does not explain the pattern of our data. We suggest that understanding of inference as a source of knowledge is genuinely difficult and requires a further conceptual advance. So far, our research acted as corrective to the existing literature. But, we also extended Sodian and Wimmer's (1987) and Keenan, Ruffman and Olson's (1994) findings in suggesting that understanding of inference as a source of knowledge rapidly reaches explicit awareness.

References

Flavell, J. H. (1988). The development of children's knowledge about the mind: From cognitive connections to mental representations. In J. W. Astington, P. Harris, & D. R. Olson (Eds.), *Developing theories of mind* (pp. 244-267). Cambridge, MA: MIT Press.

Halliday, M. S. (1977). Behavioural inference in young children. *Journal of Experimental Child Psychology, 23,* 378-390.

Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science.* Cambridge, MA: MIT Press.

Keenan, T., Ruffman T., & Olson, D. R. (1994). When do children begin to understand logical inference as a source of knowledge? *Cognitive Development, 9,* 331-353.

Lewis, C., Freeman, N. H., Kyriakidou, C., Maridaki-Kassotaki, K., & Berridge, D. (1996). Social influences on false belief access: Specific sibling influences or general apprenticeship? *Child Development, 67,* 2930-2947.

Pears, R., & Bryant, P. (1990). Transitive inferences by young children about spatial position. *British Journal of Psychology, 81,* 497-510.

Perner, J. (1991). *Understanding the representational mind.* Cambridge, MA: MIT Press.

Perner, J., Ruffman, T., & Leekam, S. R. (1994). Theory of mind is contagious: You can get it from your sibs. *Child Development, 65,* 1228-1238.

Sodian, B., & Wimmer, H. (1987). Children's understanding of inference as a source of knowledge. *Child Development, 58,* 424-433.

Wellman, H. M. (1988). First steps in the child's theorising about the mind. In J. W. Astington, P. Harris, & D. R. Olson (Eds.), *Developing theories of mind* (pp. 64-92). Cambridge, MA: MIT Press.

Wellman, H. M. (1990). *The child's theory of mind.* Cambridge, MA: MIT Press.

Wimmer, H., Hogrefe, G. J., & Perner, J. (1988). Children's understanding of informational access as a source of knowledge. *Child Development, 59,* 386-396.

Wimmer, H., Hogrefe, G. J., & Sodian, B. (1988). A second stage in children's conception of mental life: Understanding informational accesses as origins of knowledge and belief. In J. W. Astington, P. Harris, & D. R. Olson (Eds.), *Developing theories of mind* (pp. 173-192). Cambridge, MA: MIT Press.