ELSEVIER

COGNITION

Cognition 66 (1998) 1–31

Solving belief problems: toward a task analysis

Daniel Roth[a], Alan M. Leslie[b,*,1]

[a]*Department of Psychology, Tel Aviv University, Tel Aviv, Israel*
[b]*Department of Psychology, Center for Cognitive Science, Rutgers University, Busch Campus, Piscataway, NJ 08855-1179, USA*

Received 25 July 1994; accepted 25 November 1997

Abstract

Solving belief problems develops as a skill in normal children during the preschool years. To understand this process of development, it is necessary to provide an analysis of the tasks used to test preschool 'theory of mind' skills. This analysis should allow us to relate the structure of a given task to the underlying cognitive mechanisms that the task engages. In two experiments, we find that 3-year-old children show a pattern of success and failure on belief tasks that is not consistent with 'conceptual deficit' accounts. Young children possess the concept, BELIEF, but have certain characteristic difficulties with correctly calculating the contents of beliefs. In childhood autism, by contrast, the mechanisms that in normal development bestow conceptual competence in this domain are impaired. In the first experiment, parallel task structures are used to show that 3-year-olds are no better at predicting behavior from a partially true belief than they are at predicting behavior from an entirely false belief. We develop specific proposals about task structural factors that either facilitate or hinder success in belief-content calculation. These proposals are supported in a second experiment. We compare two false-belief tasks, one of which has helpful structural factors, the other of which has hampering factors, with a third task which exemplifies a hampering task structure but without any theory of mind content. We compare 3- and 4-year-olds' patterns of performance with that of autistic children. Each of the three groups shows a distinct performance profile across the three tasks, as predicted for each case by our model. Innate attentional mechanisms provide the conceptual foundations for 'theory of mind' but must be supplemented by a robust executive process that allows false beliefs to achieve 'conceptual pop-out.' Our approach has general implications for the study of conceptual development. © 1998 Elsevier Science B.V.

Keywords: Theory of mind; Childhood autism; Conceptual development

* Corresponding author. Tel.: +1 732 4456152; fax: +1 732 4456280; e-mail: aleslie@ruccs.rutgers.edu

[1] It is with sorrow that I note here the death of Daniel Roth, the first author. The work presented here was carried out in partial fulfillment of the Ph.D. degree awarded posthumously to Dr. Roth by Tel Aviv University. We began work together on writing this paper, in high spirits and with great optimism, but sadly I have had to finish the task alone. The responsibility for errors and omissions is therefore entirely mine.

PII S0010-0277(98)00005-5

1. Introduction

The discovery of the false-belief problem (Wimmer and Perner, 1983; Baron-Cohen et al., 1985) has turned out to be a landmark in the study of commonsense ‘theory of mind’. In a standard task[2], Sally puts an object into a basket and then departs. While she is away, Ann removes the object from the basket and places it in a box. The child is then asked either to predict Sally’s behavior on her return (‘where will Sally look for the object?’) or is asked directly about Sally’s belief (‘where does Sally think the object is?’). In either case, the results are the same. From around the fourth birthday onwards, most normal children can infer that Sally will still believe the object to be where she herself had put it, while younger children fail to infer the content of Sally’s belief correctly. Most autistic children, even if they have mental ages well in excess of 4 years, fail this task, while more severely retarded children (e.g. Down’s syndrome) can solve the problem with ease. Since the publication of these seminal findings, the patterns of normal and abnormal development they revealed have attracted a great deal of interest (e.g. Astington et al., 1988; Butterworth et al., 1991; Whiten, 1991; Baron-Cohen et al., 1993).

However, a decade later there has been little progress in answering the key questions these discoveries raise. Many current views assume that the young child uses general intelligence to discover a succession of ‘theories’ concerning mental states, apparently much as a mature scientist might (Gopnik and Meltzoff, 1997). To date, ‘theory-theories’ have been content simply to describe what the child does or does not know at different ages without attempting to say how any particular concept is discovered or how any particular problem is solved.

How is the concept, BELIEF, acquired? How are false-belief problems solved? The challenge to traditional theories of concept acquisition is great indeed because theory of mind concepts are deeply and ineluctably abstract. It may be that an account relying upon general induction will yet be forthcoming. But we doubt it. In the meantime, we are prepared, concept by concept, to explore the possibility that a given abstract concept might be innate, and if so, how it could be so. In the process, we hope to learn about the general nature of concepts, however the possibilities turn out.

In trying to understand the results with which we started, there is limited usefulness in simply *describing* ‘what the child knows’ or does not ‘know’ at different ages or in different clinical groups. Although precisely this sort of aim is commonplace in developmental psychology, it assumes that ‘the child’ can be treated as a unitary entity and that ‘knowing’ can be treated as a unitary property of that entity. But cognitively speaking, neither ‘the child’ nor ‘knowing’ is indivisible. Uncovering underlying cognitive architecture, the various processing mechanisms and representational systems that both produce and undergo development, is, we believe, critical for understanding conceptual development and its abnormalities (Leslie and Roth, 1993).

[2]Throughout, we shall use the term ‘standard’ descriptively to denote tasks that 4-year-olds typically pass and 3-year-olds typically fail. We do not imply that standard tasks are somehow more important than non-standard tasks.

Though admittedly still sketchy, the best specified model of underlying processing mechanisms in this domain is the ToMM theory proposed and developed by Leslie and colleagues (e.g. Baron-Cohen et al., 1985; Leslie, 1987, 1988a,b, 1991, 1992, 1994a,b; Frith, 1989; Baron-Cohen, 1990, 1995; Leslie and Frith, 1990; Frith et al., 1991; Roth and Leslie, 1991; Charman and Baron-Cohen, 1992; Leslie and Thaiss, 1992; Leslie and Roth, 1993; Happé, 1995; Leslie and German, 1995;). The ToMM theory brings both normal and abnormal patterns of development within a single neuropsychological explanatory framework. This allows us to formulate more powerful empirical methods than is possible if one studies normal and abnormal patterns in isolation or if one treats abnormalities as having a purely clinical interest. The ToMM theory also allows us to draw a principled distinction in this domain between competence and performance. The specific conceptual competence required for understanding a false-belief problem is associated with the system of (meta)representation ToMM makes available. The performance requirements for successfully solving false-belief problems – for successfully calculating the contents of false beliefs – are not entirely satisfied by ToMM. We have hypothesized that certain critical performance demands are not specific to false-belief problems but are demands made more generally (Leslie and Thaiss, 1992; Leslie and Roth, 1993). These more general demands are modeled within our framework by a distinct but related mechanism. We have dubbed this mechanism the 'selection processor' (SP). Leslie and Polizzi (in press) offer some detail on how SP co-operates with ToMM to solve false-belief problems.

The ToMM-SP framework illustrates how a highly specialized and a less specialized mechanism might work together to produce cognitive development. The less specialized mechanism leverages the limited but crucial powers of the more specialized mechanism into a system with greater scope.

1.1. Competence

On the competence side, according to the ToMM theory, the child is endowed with a set of basic attitude concepts, such as, BELIEVING, DESIRING and PRETENDING. These concepts form the core domain-specific basis for acquiring a theory of mind and are deployed within a specialized kind of data structure called the 'metarepresentation' or 'M-representation' (Leslie, 1987; Leslie and Thaiss, 1992; Leslie and Roth, 1993). The role of ToMM and the metarepresentation might be understood by analogy with a functionally specialized component of the visual system that, say, employed a basic vocabulary of shapes within a representational mechanism that could learn to recognize visual objects. Another, again partial, analogy is provided by language acquisition mechanisms that deploy universal grammatical representations to learn a natural language – mechanisms that have been dubbed the 'language instinct' (Pinker, 1994). Likewise, we might dub ToMM the 'theory of mind instinct'.

Within the M-representational data structure, attitudes are represented by three-place relational terms which we call 'informational relations.' The three argument slots of the relation specify, respectively, an 'agent', an informational 'content' (using a 'decoupled' representation), and an 'anchor' or aspect of the real world. For example: mother BELIEVES/DESIRES/PRETENDS (that) 'it is in the basket' (is true of) the marble.

In the above example, the agent (mother) is represented as holding one of the attitudes (believing, desiring, or pretending) toward the truth of the content ('it is in the basket') in regard to some particular marble. According to the model, ToMM has the job of providing such agent-centered interpretations of behavior and carries out this job by constructing M-representations.

1.2. Performance and conceptual 'pop-out'

Performance factors are often considered less important than competence. In part this reflects the context in which the distinction was introduced into cognitive science by Chomsky (1965). Accurately determining the nature of competence, often entails stripping away performance factors that obscure the underlying competence.

However, performance factors are not always 'nuisance' factors; indeed, in many theoretical contexts, their importance is fundamental. For example, in a visual texture with multiple homogenous elements, e.g. red crosses, a single element that differs from the rest in a single feature, e.g. a green cross, will be noticed immediately and without effort – it will 'pop-out' (Treisman and Gelade, 1980; Treisman, 1988; Julesz, 1991). However, if the background texture consists of conjunctions of texture elements, e.g. red crosses and green circles, and the odd-one-out element differs only in its conjunction, e.g. a red circle, then it will not be noticed immediately. To spot this element will require 'scrutiny' (Julesz, 1994) – effortful and serial attention to each of the elements of the array. The difference between pop out and scrutiny no doubt illustrates a 'performance factor' with respect to the basic competence to discriminate texture gradients. But this factor marks a major architectural distinction with great theoretical import – the distinction between early, pre-attentive, bottom up, parallel visual processing and late, attentive, serial, top down visual processing.

Our point is to suggest that understanding performance factors in conceptual development can in itself be theoretically important. In fact, the example above of perceptual pop-out might suggest a way of thinking about the development that occurs in regard to false-belief problems between ages 3 and 4 years. For the 4-year-old, a person's false belief 'pops out,' even in a standard task. For the 3-year-old, false beliefs rarely pop-out and, whenever they can be recognized, it is only through an effortful scrutiny in helpful circumstances. Given that learning about false beliefs (and allied matters such as surprise, lying, double bluff, and so forth) must remain hard so long as scrutiny is required, it is important to ask, what change in the cognitive system allows conceptual pop-out for false beliefs?

Roth (1993) (see also Leslie and Thaiss, 1992; Leslie and Roth, 1993; Leslie, 1994a) suggested that ToMM has limited powers because of its modular nature. These limits show up in tasks in which the correct solution requires the child to disregard salient information (e.g. current reality) and resist the default belief

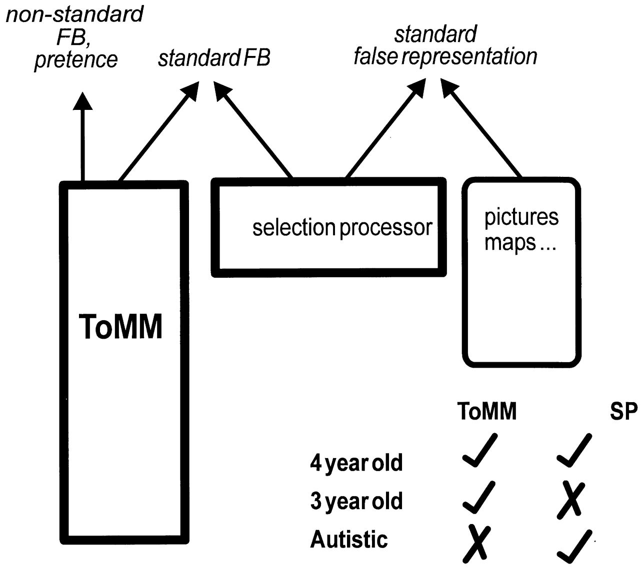

Fig. 1. The ToMM-SP model of normal and abnormal development (after Leslie and Thaiss, 1992). Four-year-olds have an intact ToMM and a well-functioning SP and can solve all the tasks shown. Three-year-olds have an intact ToMM but a poorly functioning SP which restricts some of their task performances. Autistic children show a third pattern in which SP but not ToMM functions well. This gives rise to their characteristic ‘theory of mind’ impairment.

assumption (viz. that belief contents are true). Some false-belief tasks, such as the Sally and Ann scenario, make such demands. To infer the correct content for Sally’s belief, the child must resist the temptation to select the currently baited box in favor of the empty basket as the location for the marble in Sally’s belief. The need to resist a potent response and the spatial reversal of the position of the object in Sally’s belief are reminiscent of Diamond’s analysis of the A$\overline{\text{B}}$ error in infants (Diamond, 1988; cf. Butterworth, 1993). The suggestion is that ToMM requires the help of an ‘executive’ process to solve certain false-belief problems. We dub this co-operating mechanism the ‘selection processor’ (SP) and portray its relationship with ToMM in Fig. 1.

Others (notably Russell et al., 1991) have suggested a role for executive processing in theory of mind tasks. We have two major points of disagreement with Russell. We do not believe that executive functioning is the source of the theory of mind impairment in autism; we have discussed this at length in Leslie and Roth (1993). The second difference concerns what creates the salience-related difficulty for 3-year-olds. Russell hypothesizes that, ‘physical knowledge is more salient than mental knowledge (for 3-year-olds) so that... where the two are in competition the former wins out’ (Russell et al., 1991; see also Saltmarsh et al., 1995; for a similar view). The SP hypothesis does not locate the competition as between physical and mental state knowledge but between two possible belief contents, one of which represents physical reality, the other not. Despite these disagreements, we do feel the need, along with Russell, for ‘an explanatory cognitive theory of how the acquisition described by the ‘theory of mind theory’ proceeds’ (Russell et al., 1991).

The SP allows the child to select the correct belief content in a false-belief task and to resist the default assumption that belief contents are true. In this view, solving standard false-belief tasks requires the co-operation of ToMM and SP; thus the child needs to have both components well developed. One reason to suppose that SP is a general mechanism is evidence from other tasks that have a similar problem structure to that of standard false-belief tasks. For example, the out-of-date photograph tasks of Zaitchik (1990) are highly similar to false-belief tasks in their task structure, and 3-year-olds also fail these tasks. In our model (Fig. 1), solving these other tasks also requires the availability of SP (but not of ToMM) together with knowledge concerning photographs[3].

According to the ToMM-SP model, the problems 3-year-olds have with standard false-belief tasks (and with ‘standard’ false photographs tasks) reflect performance limitations associated with SP. These limitations lead 3-year-olds to miscalculate the content of beliefs when the content is false. By contrast, autistic children’s failure on false-belief tasks reflects competence limitations stemming from an impaired ToMM rather than performance limitations associated with SP. In this regard 3-year-old normal children and older autistic children are mirror-images of each other (Leslie and Thaiss, 1992; Leslie and Roth, 1993). These aspects of the model are supported:

1. by evidence that autistic children perform near ceiling on false photographs, drawings and maps tasks while failing standard false-belief tasks (Leekam and Perner, 1991; Charman and Baron-Cohen, 1992, 1995; Leslie and Thaiss, 1992), despite the fact that the false-belief tasks and photographs tasks share the same general problem solving structure; and
2. by evidence that 3-year-old children tend to fail both standard false-belief and photographs tasks (Zaitchik, 1990) while passing other false-belief problems that do not share the critical features of problem solving structure that stress SP (e.g. Wellman and Bartsch, 1988; Mitchell and Lacohée, 1991; Roth and Leslie, 1991; Zaitchik, 1991; see discussion in Leslie and Thaiss, 1992; Freeman and Lacohée, 1995; Saltmarsh et al., 1995; and see Section 1.3).

1.3. A competing view

An alternative view has been argued for by Perner (1988,1991), (see also Flavell, 1988; Zaitchik, 1990; Wellman, 1990; Gopnik, 1993; for related views). According to Perner, the child lacks the concept BELIEF until 4 years of age. The younger child’s

[3]Another quite different possibility is that selection processing is done by ToMM ‘on-board’. Then other structurally similar tasks that demand it may have selection processing supplied by different dedicated circuitry. Let us suppose that prefrontal cortex supplies all of the distinct SP-performing circuits. If prefrontal cortex has a single maturational rate, then this will be enough to give rise to *cognitively coincidental* developmental timing in success among the different tasks. This will mimic a single general SP. Our minds remain open to such possibilities.

failure on false-belief tasks is seen as the result of an inappropriate ‘theory of representation’. Children must first develop an adequate ‘theory of representation’ and then apply this theory to their early understanding of mental states, specifically, beliefs. Prior to this, the child is unable to conceptualize *mis*representation and is thus unable to conceive of a belief that is false. The critical piece of evidence supporting this view is taken to be young children’s failure on standard false-belief tasks.

Elsewhere (Leslie and Thaiss, 1992; Leslie et al., 1993; Leslie and Roth, 1993; Leslie, 1994a), we have criticized the representational theory of mind (RTM) view, arguing that its crucial assumptions are implausible and lack empirical support. What evidence is there that the 4-year-old employs anything other than propositional attitude (PA) concepts in understanding false belief, viz., a concept of representation? Given that false-belief tasks can be solved using *only* the PA notion, BELIEVING (THAT), success on false-belief tasks cannot itself support the further assumption that the child has a representational theory of what PA’s are. What would relevant evidence look like? The critical difference between mental states as representations and mental states as propositional attitudes lies in how the states are individuated. As PA’s, mental states are individuated solely by content; as representations, beliefs are individuated by *both* form *and* content. Because false-belief tasks require only that the beliefs involved be individuated by content (and not by form), performance on these tasks does not choose between the two views.

Far from there being any evidence that preschoolers distinguish beliefs with identical contents but different forms, there is no compelling evidence that preschoolers conceive of beliefs as having any particular form at all. As Leslie and Thaiss (1992) show, belief-as-picture-in-the-head – the most plausible preschool version of RTM – can be ruled out empirically. It is easy to see why belief-as-picture-in-the-head is not the natural design for preschool ‘theory of mind.’ Given that two people, Mary and John, might have in mind exactly the same ‘representation’ (e.g. an image of it raining outside), how is the child to understand that? Mary might PRETEND it was raining while John BELIEVED it. This difference between Mary and John is critical for predicting their behavior. Yet the difference between them cannot be understood in terms of a difference in the representations Mary and John have in mind because it is the same representation in both cases (the image of it raining). Instead, the difference between Mary and John resides in the different attitude – pretending versus believing – that each takes toward the truth of what their ‘representations’ mean. Thus, even if the 4-year-old did conceptualize belief as representation, she must still employ the attitude concept, BELIEVES, in order to understand which mental state is involved and therefore what the impact on behavior will be. It is the PA concept that does all the work; it is therefore the concept BELIEVES rather than MENTAL PICTURE that the child needs to understand. Leslie (1994a) develops these points further.

More recently, in light of such arguments, Perner (1995) has retreated from his original general version of understanding representation. Rather than the child having to master a *general* theory of representation, Perner (1995) now requires the child to employ a theory of representation that is specific to understanding beliefs. In characterizing this theory, Perner draws upon Fodor's explication of the theoretical foundations of cognitive science. Fodor (1976, 1981) argues that a propositional attitude, such as the *belief that p*, can be understood as a *computational* relation between an organism and a mental representation expressing the proposition *p*. Fodor's account is intended as a scientific account of what propositional attitudes really are, much as a physicist might give an account of what physical objects really are by citing the atomic theory of matter. Nevertheless, Perner borrows the account and attributes knowledge of it to the child – with one modification: instead of the child conceptualizing the notion COMPUTATIONAL RELATION, Perner claims that the preschooler uses the concept SEMANTICALLY EVALUATES. According to Perner (1995), the child understands the metarepresentation, *John believes that p*, (in the case that *p* is false) by constructing the meta-metarepresentation, *John semantically evaluates a mental representation expressing the proposition that p*.

In order to discover and apply the above version of RTM, the child must acquire the following concepts: SEMANTIC, EVALUATES, MENTAL, REPRESENTATION, EXPRESSES, and PROPOSITION. None of these concepts is obviously less abstract or less obscure than the concept BELIEVES and it is certainly puzzling how the child might acquire these six notions or indeed any of them. But even if the child *did* discover this complex theory, it would still not provide the child with a concept of belief. Fodor's original formulation is designed to apply equally to all PA's; likewise, Perner's 'meta-metarepresentation' will tell the child only about propositional attitudes *in general*, applying to desires and pretends equally as it applies to beliefs. It will even apply just as well to 'preliefs', the undifferentiated pretend-belief states that Perner et al. (1994) claim 3-year-olds attribute to others. What it will not do is tell the child specifically what a 'belief' is, as opposed to these other, and earlier developing, notions.

1.4. Three-year-old's difficulty with false belief: competence or performance?

Perner's latest RTM theory is unnecessarily complex, yet still does not account for the specifics of the belief concept. However, there still remains the important question of how to understand the difference between 3- and 4-year-old performance on false-belief problems.

In some versions of the conceptual deficit view (e.g. Flavell, 1988; Wellman, 1990; Gopnik and Wellman, 1995), the *falseness* of the belief plays the critical role in creating difficulty for the child. In the case of understanding a true belief, the belief is a 'copy' of reality and there is no need for understanding misrepresentation. Likewise, a case in which an agent knows something true about a situation but does not know the whole truth (partial ignorance) will be easier to understand because this does not require representing a misrepresentation but simply the lack of 'connection' between an agent and one aspect of the situation.

In Perner's view of the conceptual deficit, the 3-year-old's lack of a representational theory of belief rules out an understanding of misrepresentation but also impacts the child's understanding of knowledge (true belief). With regard to true belief, the 3-year-old child is a 'situation theorist', that is, he conceives of people having relations to situations (rather than representations of situations). Thus, one person can desire one situation, while another desires a second situation. Although Perner (1991) is less explicit about this, presumably the situation theorist 3-year-old can conceive of knowledge in a similar way: as a ‘mentalistic’ relation between a person and a situation. This would allow the 3-year-old to grasp that, in a case of partial ignorance, a person can be related to one situation (object 1 in *x*) but not related to another situation (object 2 in *y*), if for example, the person saw object 1 in *x* but not object 2 in *y*. And indeed, Perner argues that ‘3-year-old children know something about the fact that knowledge and perceptual activities, especially seeing, go together’ (Perner, 1991). However, for the 3-year-old the relation between *seeing* and *knowing* is merely a ‘rule’ in a theory of behavior. The young child does not appreciate ‘the real importance of informational access’. For this reason, informational access as a determinant of a person’s knowledge is often less salient to 3-year-olds than, for example, the person’s age. However, given their ability to grasp mentalistic relations to different situations and their ‘see-know’ behavioral rule, and provided that informational access can be made sufficiently salient, Perner should expect better success among 3-year-olds on partial ignorance tasks than on false-belief tasks (see also Hogrefe et al., 1986; Pratt and Bryant, 1990).

The ToMM-SP theory takes a subtly different view. It does not claim that partial true-belief situations are *inherently* easier for the child to deal with than false beliefs. Instead, it depends on how easy it is for the child to accurately calculate belief content. This in turn depends upon the details of the task structure in relation to stressing SP (and, of course, on the strength of SP and the nature of its cooperation with ToMM). If a given partial true-belief task stresses SP to a degree equal to the stress produced by a given false-belief task, then the 3-year-old should perform just as well or just as badly on both.

Both sets of views can agree that it is an easier problem simply to register the fact that a character does or does not know something than it is to attribute a specific belief. Where the theories differ is in the account each gives of what makes this easier. The RTM view says that it is because understanding that someone does or does not know something does not require a ‘theory of representation,’ just a grasp of whether an agent is or is not ‘connected’ to a situation. For the ToMM-SP model, difficulty is determined by the difficulty of calculating belief contents. The easiest content to ‘calculate’ is to assume that someone knows everything you do. The next easiest content to calculate is a minimally specified content, such as ‘he knows something about this/nothing about this’. Under many circumstances, such minimally specified content will allow only a gross prediction of behavior. But the more specific the content that has to be specified to allow more accurate prediction of behavior, the more calculation difficulty will increase, other things being equal.

We tested between the above theories by using a scenario first developed by Leslie and Frith (1988) in a study of autistic children. Leslie and Frith produced a partial true-belief task by adapting their ‘Sally and Ann’ false-belief scenario (Baron-Cohen et al., 1985). The two tasks were minimally different. In the false-belief task, a live actor hid an object in location A. The actor then left the room. With the actor gone, the experimenter moved the object to location B. The child was then asked to predict where the actor would look for the object upon returning. In the partial true-belief task, the actor watches while the experimenter hides an object under location A, and then leaves the room. With the actor gone, the experimenter produces another object identical to the first. This is then hidden in location B. The child is asked to predict where the actor will look for an object upon returning.

In both the above tasks, we can ask the child if the actor *knows* there is an object in location B. As noted above, both theories predict that this question will be relatively easy for young children. On the prediction of behavior question the ToMM-SP model says that there will be little or no difference between the partial true-belief scenario and the false-belief scenario. It will not be easier for the children to correctly predict an actor's behavior when two objects simultaneously occupy two locations, but only one location is known about, as compared with the case where a single object occupies two locations sequentially but is seen in only one of them. Both cases pose 'selection' difficulties. The RTM model, however, predicts it will be easier to predict behavior on the basis of a belief which is true plus a 'does not know about location B', because neither requires understanding misrepresentation.

We examined a (partial) true-belief task structure that closely parallels the structure of a standard false-belief task. Because Leslie and Frith (1988) did not test normal children, we re-ran the Leslie and Frith study using three age groups of 3-year-olds. This allowed us to gauge whether children pass behavior prediction in partial true-belief scenarios before they pass in a task-equated false-belief scenario.

2. Experiment 1

2.1. Method

2.1.1. Subjects

Forty-seven 3-year-olds in three age groups were tested. In the 'younger' age group there were 16 children aged between 3 years and 3 years 4 months (mean age 3 years 2 months). In the 'middle' age group there were 16 children aged between 3 years 5 months and 3 years 8 months (mean age 3 years 6 months). In the 'older' age group there were 15 children aged between 3 years 9 months and 4 years (mean age 3 years 10 months). In addition to these a further 14 children were seen but were rejected, in approximately equal numbers from each group, for failing control questions or because they were uncooperative and did not complete the study.

2.1.2. Procedure

Three tasks were carried out in the following fixed order: line of sight task, partial true-belief task, and false-belief task. We used this order so that our results could be compared with those of Leslie and Frith (1988). We dropped Leslie and Frith's 'memory for position' task since we thought it served little purpose with normal children and in any case memory for position was tested by control questions in the last two tasks. Different materials were used in different tasks to minimize interference.

2.1.2.1. Line of sight task. The child was introduced to a doll and a screen. The screen was placed perpendicular to the child and experimenter. The doll could be placed on either side of the screen, always visible to the child. A plastic counter was introduced and placed on one side of the board. The child was then asked if the doll could see the counter. Depending upon the relative positions of the doll and counter the correct answer was either yes or no. Five trials were given varying the correctness of the yes or no answer. If a child failed on any trial a further three trials were given, making a maximum of eight.

2.1.2.2. Partial true-belief task. Three distinctive colored paper hats to be used as hiding places were introduced to the child. An actor then walked past the open door to the room, apparently fortuitously, and was recruited to help in ‘our game.’ The door was closed and the experimenter produced a coin and asked the child to name it. Whatever name the child gave it was then used to refer to the coin throughout. The experimenter then proceeded to hide the coin slowly and deliberately under one of the hats. The experimenter asked the child and the actor in turn, ‘Can you see me hiding this (coin) under here?’ The actor made a point of agreeing that she could. Having completed the hiding, ensuring the child attended both to the place of hiding and to the fact that the actor could see this, the experimenter then asked the actor to leave the room. While the actor was outside with the door closed again, the experimenter drew the child’s attention to the fact that the actor was no longer in the room and could not see what they were doing. The experimenter produced another identical coin and showed this to the child saying, ‘Look, I have another (coin).’ This was then given to the child with the request, ‘Can you put this somewhere different?’ (control question 1). All children successfully chose a hiding place different from the previous place, indirectly demonstrating memory for the original location. The child was then asked, ‘Where did (the actor) see me hide a (coin)?’ (control question 2). The response here had to be correct or else the task was repeated. A second failure on a control question led to the child being rejected. After the child’s correct response, the experimenter pointed to the place where the child had hidden the second coin and asked, ‘Does (the actor) know that there is a (coin) under here?’ (Know question). Finally, the child was asked, ‘When (the actor) comes back in, where will she look for a coin?’ (Prediction question). Here the child could point to any or all of the three locations. Pointing to the location, and only that location, where the actor had seen the coin being hidden was considered correct.

2.1.2.3. False belief. The three hats were removed and replaced by a toy basket with a cloth cover, a red purse and a small brown box. The experimenter produced a marble, showed it to the child agreeing with him/her what it should be called, then gave it to the actor saying, ‘Now, will you hide this somewhere?’ The actor said, ‘OK. Look, I’m going to hide (the marble) in the basket’ and did so. The actor was then asked to leave the room again. The experimenter again drew the child’s attention to the fact that the actor was no longer in the room and could not see what we were doing. Then the experimenter asked, ‘Where did (the actor) hide (the marble)?' After the correct response, the experimenter removed the marble from the basket ensuring the child was attending and hid it in the brown box. The child was then asked the following questions: 'Does (the actor) know the marble is in here?' (know question), pointing to the new location; 'When (the actor) comes back in, where will she look for the marble?' (Prediction question); 'Where did (the actor) put (the marble) in the beginning?' (control question 1); and 'Where is (the marble) really?' (control question 2). Five children who failed one or both of these control questions were rejected.

2.2. Results

Our criterion for passing the line of sight task was 5 out of 5 correct or 7 out of 8 correct (i.e. $P < 0.05$). All the children apart from two subjects in the young age group passed this strict criterion.

The percentages of children passing the know question in the partial true-belief task and in the false-belief task is shown by age group in Fig. 2. Children's performance on the know question improved dramatically between 3 years 2 months and 3 years 6 months. There were significant differences between the younger group and the rest, both in the partial true-belief task (Fisher's Exact, $P < 0.001$) and in the false-belief task (Upton's $\chi^2 = 15.5$, $P < 0.001$). Performance on this question across the two tasks was very similar in all age groups.

The percentages of children passing the prediction question in the partial true-belief task and in the false-belief task is shown by age group in Fig. 3. There was a significant effect of age only in the false-belief task (older versus rest, Fisher's Exact, $P < 0.03$). Children's performance on prediction shows improvement in the months leading up to the fourth birthday. This improvement was slightly

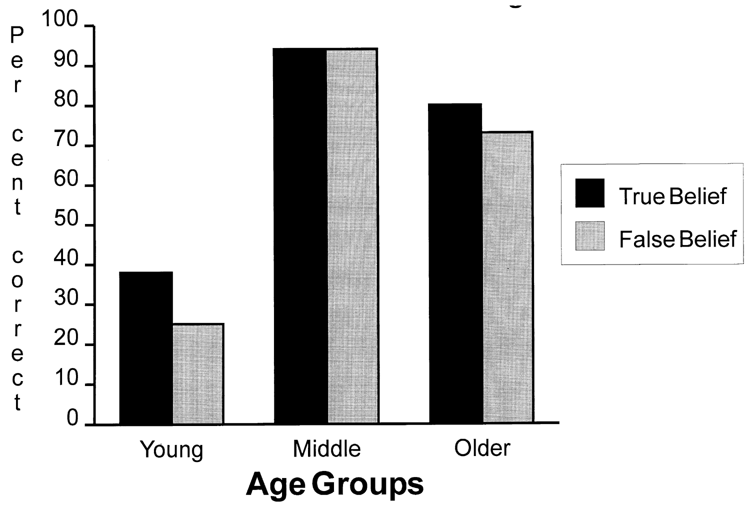

Fig. 2. Percentages of 3-year-olds in the three age groups passing the 'know' question in the partial true-belief and false-belief tasks (experiment 1).

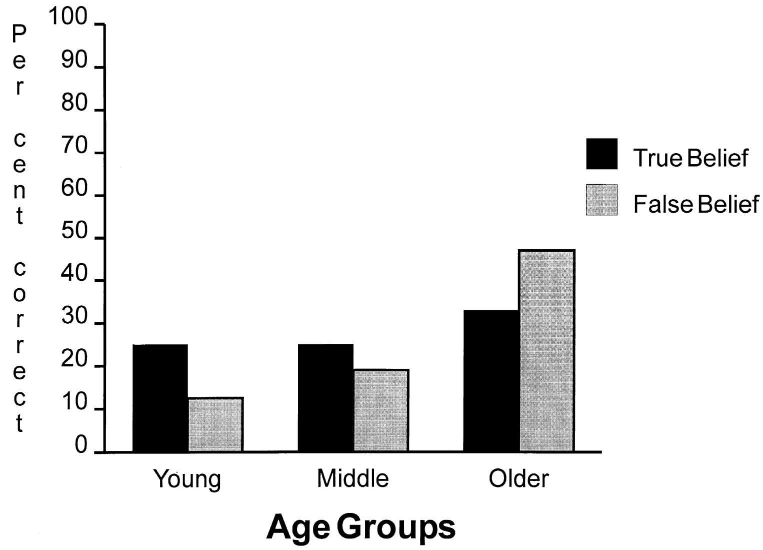

Fig. 3. Percentages of 3-year-olds in the three age groups passing the ‘prediction’ question in the partial true-belief and false-belief tasks (experiment 1).

more marked in the false-belief task but prediction in false belief was not easier than in partial true belief: in the older group, two children predicted correctly in partial true belief but not in false belief, while four children showed the opposite pattern (Binomial, $P > 0.3$).

Of the children in all age groups who failed the true belief prediction, 47% pointed to the new location (i.e. where they themselves had placed the other coin), 26.5% pointed to both old and new locations, and 26.5% to the irrelevant location or to all three locations). Of the children who failed the false belief prediction, 78% pointed to the new location.

We compared the present results with those obtained by Leslie and Frith (1988) with autistic children on the same tasks. In that previous study, 18 autistic children with a mean verbal mental age of 7 years 2 months performed similarly to the present middle aged 3-year-olds on false belief prediction and rather worse than the older group (28% passed). On true belief, the autistic children were 61% correct on know and 50% correct on prediction. On this task, then, the autistic children of Leslie and Frith (1988) showed a pattern not found in any of our 3-year-old groups. This becomes more obvious when one considers the relationship between performance on know and prediction (see Fig. 4). In the middle 3-year-old group, all but one child passed the know question whereas only 27% of these went on to predict behavior correctly. In contrast, 73% of the autistic children who passed know were also able to predict behavior correctly (Upton’s $\chi^2 = 5.21$, $P = 0.022$, two-tailed). For all 3-year-olds, 70% passed the know question; of these know passers, only 33% went on to predict behavior correctly, again in contrast to the autistic children (Fisher’s Exact, $P = 0.027$).

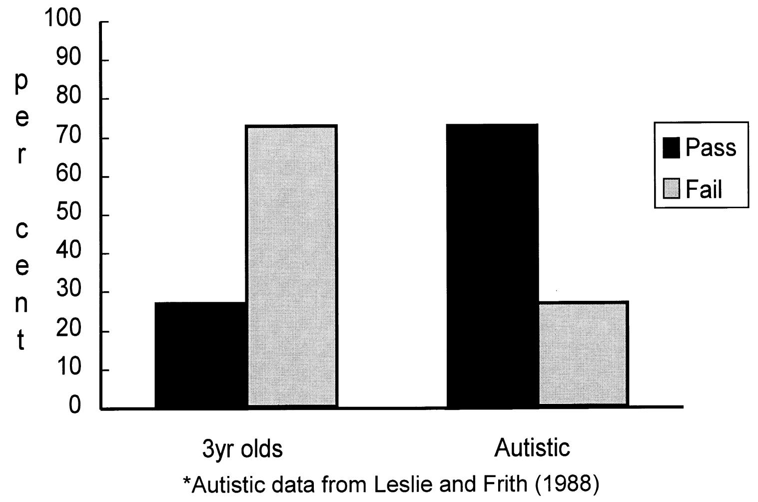

Fig. 4. Three-and-a-half-year-old and autistic children show different patterns of responding on the partial true-belief task. For autistic children, success on prediction follows success on know, but not so for 3-year-olds.

2.3. Discussion

Within the ‘standard’ task format used here, children pass knowing/not knowing questions earlier than prediction of behavior questions. There was a sharp rise in success on the knowing/not knowing question from the young to the middle aged 3-year-olds, in line with the findings of Pratt and Bryant (1990). We think it unlikely that this was the result of a ‘no’ bias in answering this question. First, only our two older normal groups tended to answer ‘no’, while the youngest group tended to answer ‘yes’. Second, as Perner (1991) has argued, when young children do not understand know questions, they tend to attribute knowledge to adults regardless of the adults’ access to information. In our study, this would have produced a ‘yes’ bias, opposite to what was found in the older two groups. Our older 3-year-olds’ appear to have understood our knowing question.

Success in prediction, however, lagged considerably behind success on the knowing question, even in the oldest 3-year-olds. Our results on this depart from those of Sullivan and Winner (1991) who found that success on ignorance and false belief came together for older 3-year-olds. However, in that study children were tested on a think (rather than prediction) question in a Smarties-type task rather that on prediction in a Sally and Ann-type task as in the present study. On prediction of behavior, we found no difference between the partial true-belief task and the false-belief task. This suggests that, when the task structures are minimally different, it is not easier for children to predict behavior from partial ignorance than from false belief. Given that our older group achieved nearly 50% success on false belief and given that, on the basis of previous studies, one can expect around 80% success from four-and-a-half-year-olds, it seems unlikely that a dramatic difference in favor of true belief over false belief will open out after the fourth birthday.

These findings are striking from the point of view of a number of current ideas. First, although the 3-year-olds from 3 years 6 months onwards show a clear understanding of the actor's 'informational access' (Wimmer et al., 1988) to one location but not to the other, this does not help them predict behavior. Likewise, although the 3-year-old understands that the actor knows about and is therefore presumably 'thinking of' one location but not the other (Perner, 1989), they do no better at predicting behavior in the true-belief scenario than in the false-belief scenario. Alternatively, if 3-year-olds use a 'copy theory' of belief (Wellman, 1990), then they will know that the actor has a mental 'copy' of the coin in location A but has no mental 'copy' of the coin in location B, and again the result should have been a successful prediction of behavior.

In a somewhat different vein, one might have expected very similar results to be obtained from Wellman and Bartsch's version of the true-belief task ('inferred belief control') and from the present version (Wellman and Bartsch, 1988). Wellman and Bartsch told 3-year-olds a short story in which there are magic markers in two locations, but a character sees them in only one of the locations. The children were asked to predict where the character would look. In Wellman and Bartsch, 88% of three-and-a-half-year-olds succeeded in predicting the behavior of the protagonist, while in the present study most failed. Given the close similarity of *conceptual content* between these two versions of true belief, the very different results are presumably due to differences in the processing demands made by the different task structures. In the Wellman and Bartsch study, the child is told directly by the story teller about the exposure history of the protagonist: Jane saw this but did not see that. This may reduce the inhibition and selection problems. In the present study, the child is given as little help as she is given in the standard false-belief version. Apparently, this leaves the child vulnerable to the standard error in the calculation of belief content.

Finally, we should also note that the present findings reflect upon Fodor's explanation of 3-year-old success in the Wellman and Bartsch scenario (Fodor, 1992). According to Fodor's model, their success in that scenario is due to the child hitting ambiguity in her prediction of the actor's behavior from the actor's desire. The actor desires the object, but, given that identical objects are in two places simultaneously, in which location will the actor look, A or B? Since the prediction is ambiguous, the child switches heuristic and, for the first time, consults belief. As soon as she does this, she succeeds. The present findings do not reflect upon Fodor's model in general because the child may not regard the action of looking in two places as ambiguous. However, the results do suggest that ambiguity was not what allowed 3-year-olds to succeed in Wellman and Bartsch's task, because exactly the same ambiguity is present in our own task and the children fail. Some other explanation is required.

2.3.1. A comparison of normal 3-year-old and autistic children

Before continuing our exploration of the ToMM-SP model, we want to consider the comparison between the present results and those Leslie and Frith (1988) obtained with autistic children. On prediction from false belief, Leslie and Frith's autistic children, with a mean verbal mental age of 7 years 2 months, performed like our middle aged 3-year-olds and slightly worse than our older 3-year-olds at 28 versus 47% passing, respectively. The autistic children even did less well than the middle 3-year-olds on the knowing question in true belief (61 versus 95%, respectively). However, we do not wish simply to underline the now familiar impairment of autistic children in such tasks but to reflect upon an intriguing pattern that emerges from a more detailed comparison. The autistic children showed a high degree of consistency between success on the knowing question and success on prediction of behavior in the partial true-belief task. Nearly three-quarters of those who were aware that Sally did *not* know there was a coin in location B correctly predicted that she would look in location A. Recall that all of the autistic children were able to pass an independent test of 'seeing and not seeing'. These patterns are consistent with the idea that the successful autistic children might employ a kind of 'behaviorist' strategy – a sort of 'SEE → KNOW → BEHAVE' rule. If Sally sees the coin in A, she will 'know about' the coin in A, and therefore be disposed to act toward A (to satisfy her desire for a coin). By the same token, if Sally does not see the coin in B, then Sally will not 'know about' or be 'associated with' B, and will not be disposed to act toward B (to satisfy her desire for a coin). This 'behaviorist' approach will not extend successfully to the false-belief situation because Sally, of course, cannot act toward the coin by going to an empty location.

Interestingly, something very similar to the above 'see → know/be associated with → behave' rule has been proposed as a characterization of normal 3-year-old's theory of mind prior to the emergence of true 'representational' conceptions (e.g. Perner, 1988, 1989, 1991). Such a proposal may well be applicable to some of the more successful autistic children. However, in the present study, normal 3-year-olds, despite showing at or near ceiling performance on both 'seeing' and 'knowing', nevertheless failed to base their predictions of behavior on this simple strategy. We believe that the reason they do not employ this simple stratagem for success is that, unlike autistic children, normal 3-year-olds follow a mentalistic instinct and attempt to calculate belief – even at the cost of failure.

2.3.2. Why is false belief difficult for 3-year-old children?

We have already discussed one answer to why 3-year-olds fail standard false-belief tasks: Beliefs must be understood as representations and children at that age do not have an adequate theory of representation. According to one version (Perner, 1988, 1991), it is the falseness of the belief that causes the trouble. Young children are unable to appreciate *mis*representations as interpretations of reality, and therefore simply use their own knowledge of reality. According to another version of this approach, the difficulty young children have with beliefs derives from an inappropriate understanding of representations as *veridical* interpretations of reality (Zaitchik, 1990). Yet other versions of the conceptual deficit view hold that young children understand beliefs only as *copies* of reality (Chandler and Boyes, 1982; Wellman, 1990; Gopnik and Wellman, 1995). None of these views are supported by the results of the first experiment.

The theory of ToMM provides an alternative to the conceptual deficit view. ToMM endows (2- and) 3-year-olds with an elementary understanding of propositional attitudes. This basic competence forms the specific basis for learning about mental states. However, this in no way guarantees that the child will be able to solve any theory of mind task we can devise. Success will depend upon the exact constellation of problem solving resources required by a given task and upon the availability of these resources at a given point in development. Certainly, we should not expect preschoolers to be able to correctly calculate the content of beliefs equally in all situations. Their conceptual competence may be ‘squeezed’ by a number of performance factors. It is plausible that, as a result of maturation, learning, practice or all of these, performance ‘squeeze’ will gradually relax over time. None of this should be particularly surprising. Indeed, on this view, the steady rise in success on both tasks in the first experiment was entirely expected. We are not required to posit a sudden or radical shift in conceptual structure.

A rigorous theory of task structure would relate success and failure to the demands made by a given task on available cognitive processes. Undoubtedly, it will be difficult to develop such a theory. However, we can try to develop studies that pull apart the demands of general problem solving from the demands of conceptual content. Sensing this imperative, Zaitchik (1990) devised what she called her ‘gizmo task’. This involved a simple mechanical device that could move between one of two states. Three-year-olds had no difficulty in following the transitions between states. However, this task was not sufficiently similar in structure to the false-belief task to show that 3-year-olds can meet its general demands. False-belief tasks do not have just a single entity that changes state but also a second parallel entity, with a strong similarity to the first which is initially tied to the first state but which then does not change. Zaitchik’s ‘gizmo task’ fails to mimic this critical aspect of task structure. Ironically, it is Zaitchik’s photographs tasks that provide a structure isomorphic to standard false-belief tasks. But, of course, as Zaitchik shows, three-year-olds fail these as well.

Notwithstanding Gopnik’s claim (Gopnik, 1993) that performance factors can be ignored, Zaitchik’s results support the idea that 3-year-olds cannot meet the general problem solving demands made by standard false-belief tasks. Autistic children who fail false-belief tasks by contrast pass the same photographs tasks. Only in the case of autistic children, then, can we rule out general problem solving limitations as the reason for their failure on false belief. We turn now to examine further the nature of the processing squeeze on 3-year-olds.

2.3.3. Probing the ToMM-SP model

To further our task analysis, we designed the following experiment probing the ToMM-SP model. We included three tasks. The first task tries to recreate part of the structure of standard false belief but without any representational content. The second aims to provide extra structure within a false-belief task in order to ‘simplify’ the belief content calculation. The third is a standard false-belief task included for reference purposes.

In the first task, dubbed the ‘screen’ task, we aim to load on selection processing outside of the context of a mental state or representation task. The child is first presented with an opaque screen. A toy basket and a box are placed in front of the screen and a marble is placed into the basket. The basket (with marble) and box are then removed behind the screen and an identical set of basket, box and marble are placed in front of the screen. The new marble is placed again in the (new) basket but is later removed to the box. The child is then asked, ‘Behind the screen, where is the ball?’ This creates a partial analogue of the ‘Sally and Ann’ standard false-belief task with the basket-box-marble set behind the screen corresponding to the information stored in the agent’s belief. The task is also partially analogous to the false photographs task. We say partial because our intuition is that the relationship between a belief or a photograph and the state of affairs it describes may have a ‘closeness’ that two merely similar states of affairs cannot have. Holding them apart may stress SP to a higher degree than we can achieve with our task. Nevertheless, though conceptually and structurally simpler and lacking strong ‘parallelism’, the ‘screen’ task may yet require selection processing to some measurable degree. If so, it will be without the involvement of either mental states or representational artifacts.

The second task, dubbed the ‘search’ task, is a non-standard false-belief scenario. The child is shown two characters both with the same desire to eat cookies. One character searches in one place and finds cookies, while the other character searches in another place and fails to find cookies. The behavior of the two characters is thus contrasted, one satisfying desire and the other failing to satisfy desire. The contrast between the two characters, the failure of desire satisfaction, and the search-behavioral clues as to the content of their respective beliefs may conspire, individually or together, to increase the potency of the false content (reducing the need for inhibition) and aid its selection. In short, such a task structure should relieve the stress on SP.

Our third task was a standard false-belief task. We tested three groups in a within-subjects design: normally developing 3-year-olds, 4-year-olds, and older autistic children. Our model predicts different patterns of performance across the groups. Four-year-olds should do well on all three tasks. Three-year-olds should do better on the non-standard search than on the other two tasks. Autistic subjects should do better on the screen than on the other two tasks. Finally, if, in normal development, selection processing is a limiting factor on standard belief problems, we should expect to see a correlation in normally developing subjects between performance on the ‘screen’ task and standard false belief but not between non-standard ‘search’ and standard false belief. This should not hold for autistic subjects, because, according to our model, SP is not a limiting factor for them.

3. Experiment 2

3.1. *Method*

3.1.1. *Subjects*

Seventy-eight children were tested on three tasks. Due to experimenter error, age data on one 3-year-old, two 4-year-olds, and three autistic children were not obtained. In addition, mental age data were collected for only 18 of the autistic children. Thirty-three 3-year-olds aged between 3 years and 3 years 11 months (mean age 3 years 6 months) from medium and medium-high socio-economic nursery schools in the Tel Aviv area were tested. Twenty-four 4-year-olds between 4 years and 5 years (mean age 4 years 5 months) from a similar background were also tested in Tel Aviv. Twenty-one autistic children, diagnosed according to standard criteria (DSM IIIR, 1987), aged between 6 years and 19 years 1 month (mean age 13 years 4 months) with verbal mental ages between 3 years 11 months and 11 years 6 months (mean 5 years 6 months) were tested in autistic schools in England. We tested the autistic children in England because suitable verbal mental age tests in Hebrew were not available.

3.1.2. Procedure

Each child was tested individually in a quiet room. The tasks were acted out using puppets and cardboard props with the experimenter sitting opposite the child. The tasks are detailed below.

3.1.2.1. Standard false-belief task. The child is presented with a doll's house made of cardboard in which there is a bed. Two characters are introduced, a bear and a white rabbit, and the following story is told: 'This is Teddy Bear. He likes to eat candy. So, now he has chocolate in his hand which he is about to eat. (Teddy bear actually has a chocolate in his hands.) But he cannot eat it now since it is very late and he has to go to bed. What will Teddy do with his chocolate when he goes to bed? He puts it under his bed. He will eat his chocolate in the morning, when he wakes up. (Teddy is seen going to bed, and the child is told that he is fast asleep.)

'But when Teddy is asleep, the white rabbit comes in. The rabbit smells something – sniff, sniff – it must be the chocolate under Teddy's bed! Yes it is: it's the chocolate white rabbit smells. Rabbit also wants some chocolate. So he takes the chocolate from under the bed and he puts it in his box. Now the rabbit leaves and goes off to sleep.

'Now, it's morning. Teddy Bear wakes up. He wants his chocolate.'

Belief question: 'Where does Teddy Bear think his chocolate is?'

Control questions: memory question: 'Before Teddy went to sleep, where did he put his chocolate?'; reality question: 'Where is the chocolate now?'

3.1.2.2. Non-standard false-belief 'search' task. The scenario used the same cardboard house in which was placed a toy cupboard and a round tin on a table. Then two characters were introduced: Yoram/Sam (a male doll) and Rivka/Judy (a female doll). We made sure the child could identify the characters before proceeding with the story:

'Yoram and Rivka are playing in the playground. Each one has a toy in their hand. But now Rivka is hungry and she wants a cookie. She goes into the kitchen to look for one. Rivka enters the kitchen and goes straight to the cupboard. She opens the cupboard door and finds some cookies. She takes one and goes outside to eat it.

But who is there outside? Who sees Rivka with her cookie? It's Yoram, her brother. Yoram also likes cookies and now he also wants one. So Yoram says (to the child): 'I'll also go to the kitchen'. 'Where are the cookies?' he says to himself. Then Yoram sees the tin on the table. So he says 'Ah, the cookies must be here in the tin!!!'. Yoram picks up the tin and opens it. But... it is empty! There is nothing in the tin! He closes the tin again and says: 'Oh dear, Rivka must have eaten all the cookies' and goes outdoors.'

Belief question: 'When Yoram/Sam went into the kitchen, where did he *think* the cookies are?'

Control question: 'Where are the cookies really?'

3.1.2.3. SP-control 'screen' task. The child is first presented with an opaque screen placed on the table. Then, a basket, a box and a marble are introduced. The marble is placed in the basket and covered, and the child is asked to indicate what the experimenter has done with it. When she answers correctly, the display (basket, box and marble) is removed behind the screen and it is made sure that the child cannot see them.

Then, another identical set of basket, box and marble are produced and the experimenter ensures that the child is aware of their similarity with the previous set. The ball is placed again in the basket. Then a puppet called 'Doggy' appears. He wants to take the basket for shopping. The basket, however, is heavy and so the puppet removes the marble from the basket to the box. The child is asked again to indicate what the puppet has done with the marble.

Test question: 'Where is the marble *behind* the screen?' (experimenter indicates, by pointing and looking, *behind* the screen).

Control question: 'Where is the marble *in front* of the screen?' (experimenter indicates, by pointing and looking, *in front* of the screen).

3.2. Results

Due to experimenter error, data was not obtained from three of the autistic children on the 'search' task. There were few failures on control questions. Out of 312 responses to control questions, only 13 were incorrect: among the 3-year-olds, there were two incorrect responses to control questions in standard false belief, two in 'search', and one in 'screen'; among the 4-year-olds, there were six incorrect responses in standard false belief, and two in 'screen'. None of the autistic children failed control questions.

The percentage of each group passing on all three tasks is shown in Fig. 5. Three different patterns of responding are evident. As expected, the 4-year-olds performed uniformly well on all three tasks. Again as expected, the 3-year-olds performed best on the non-standard false-belief 'search' task, better than on the SP-control 'screen' task (McNemar Binomial, $P = 0.048$, one-tail) and much better than on standard false belief (McNemar $\chi^2 = 16.1$, $P < 0.001$), though only 9% passing is probably below the norm for this age (see e.g. Section 2). As predicted, the autistic group showed yet another pattern, performing best on the 'screen' task. Their performance

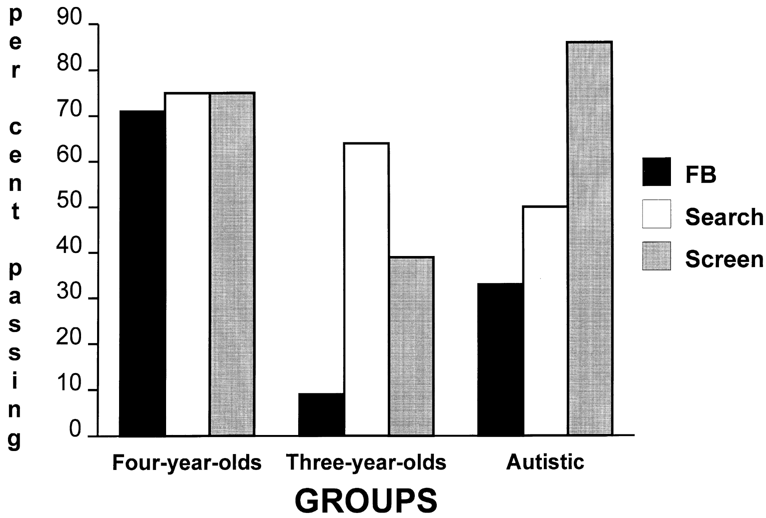

Fig. 5. Percentage of 4-year-old, 3-year-old and autistic children passing standard false-belief task, non-standard false-belief 'search' task, and 'screen' task (experiment 2).

on 'screen' is superior to that on 'search' (McNemar Binomial, $P = 0.02$), and to that on standard false belief (McNemar $\chi^2 = 7.69$, $P < 0.005$).

Due to the 3-year-old's unusually poor score on the standard false-belief task, our autistic sample out-performed the 3-year-olds, though the difference failed to reach significance (Fisher's Exact, $P = 0.064$, two-tail). The performance of the autistic children on this task was, as usual, significantly poorer than that of the 4-year-old group (Upton's $\chi^2 = 6.19$, $P < 0.01$, one-tail). Finally, if we compare performance on the non-standard 'search' false-belief task with performance on the standard false-belief task, for the autistic children there was no difference (McNemar Binomial, $P > 0.2$), while for the three-year-olds 'search' was much easier (McNemar Binomial, $P < 0.001$).

On the 'search' task, the autistic children's performance was worse than the 4-year-olds' (Upton's $\chi^2 = 2.73$, $P < 0.05$, one-tailed) and slightly but not significantly worse than the 3-year-olds'. Their performance on the 'screen' task was at least as good as the 4-year-olds' and substantially better than the 3-year-olds' (Upton's $\chi^2 = 11.05$, $P < 0.001$, one-tailed).

We examined task contingencies by counting the joint frequencies of passing and failing each pair of tasks. To look at contingencies in normal development, we combined the 3- and 4-year-old groups. Their performance on standard false-belief and SP-control screen tasks was significantly correlated ($\varphi = +0.38$, Upton's $\chi^2 = 8.01$, $P = 0.002$, one-tailed).The screen was easier than standard false belief (McNemar Binomial, $P = 0.01$). Performance on the two false-belief tasks was also correlated ($\varphi = +0.29$, Upton $\chi^2 = 4.58$, $P = 0.016$, one-tail), with 'search' easier than standard (McNemar Binomial, $P < 0.001$, one-tailed). As expected, the SP-control screen task was not correlated with the non-standard 'search' false-belief task ($\varphi = 0.03$, n.s.). The autistic children showed a different pattern of contingencies. First, as predicted, the performance on SP-control screen and standard false-belief tasks was not related ($\varphi = 0.0$), nor was the performance on screen and non-standard false belief ($\varphi = 0.1$, n.s.). The two false-belief tasks were strongly correlated ($\varphi = 0.8$, $P = 0.001$, Fisher Exact) but without differing in difficulty. See Table 1 for a summary.

3.3. Discussion

The ToMM-SP model made three main predictions:

1. there will be different specific patterns of performance in each of the three groups;
2. non-standard false belief will be easier than standard false belief for normally developing children, but no different for autistic children;
3. the Screen control task will correlate with standard false belief but not with non-standard false belief for normally developing children, and with neither for autistic children.

Each of these predictions was confirmed and are discussed below.

3.3.1. Profiles

We obtained different patterns of performance in each of the groups. This extends previous findings (Roth and Leslie, 1991) that autistic children are different on theory of mind tasks not only from 4-year-olds but also from 3-year-olds. Four-year-olds showed a consistently high level of performance across all three tasks. The 3-year-olds performed poorly on the standard false-belief task but substantially better on the non-standard ‘search’ false-belief task. Their performance on the ‘screen’ control task was intermediate and poor relative to the other groups. The autistic group performed rather poorly on both false-belief tasks but achieved a high level of performance on the ‘screen’ task.

These patterns are in line with predictions based on the ToMM-SP model of development. According to this model, standard false-belief tasks require the cooperative operation of ToMM and SP. The non-standard false-belief ‘search’ task

Table 1
Experiment 2: task correlations (4-fold point correlation, φ)

	Standard FB	Non-standard FB	Screen
Normally developing children 3–4 years old			
Standard FB			
Non-standard FB	0.29*		
Screen	0.38**	0.03 (n.s.)	
Autistic children			
Standard FB			
Non-standard FB	0.80***		
Screen	0.0	0.1 (n.s.)	

FB, false-belief. *$P < 0.025$, **$P < 0.01$, ***$P < 0.001$.

requires the operation of ToMM but places little demand on SP. The ‘screens’ control task does not require ToMM but makes selection processing demands. The model predicted the uniformly high performance on all three tasks found in the 4-year-olds because only this group is said to have both an intact ToMM and a well functioning SP.

3.3.2. Belief tasks

According to our model, 3-year-olds will do better on tasks which load purely on ToMM than on tasks which load on ToMM and SP together. The standard and non-standard false-belief tasks have task structures that load differentially on SP. In ‘search’, there are contrasting protagonists who exhibit different behaviors in pursuit of a common desire, one successful, the other unsuccessful. This contrast may increase the likelihood of inhibiting the default assumption that both characters share the same (true) belief. At the same time, the anomalous search behavior of one of the protagonists, together with overtly verbalizing his mistaken assumption, may serve to increase the salience of the false content, reduce the need for inhibition, and ‘select’ the correct content for his belief. These aspects of task structure may relieve the 3-year-old SP of the need to function efficiently or strongly. In effect, the experimenter can take over the role of SP for the child with a non-standard task.

Autistic children, by contrast, are not limited by SP in solving false-belief problems and consequently are not helped by these palliatives. For both normally developing and autistic children the two belief tasks were correlated, highly so for the autistic children who either passed or failed both. This may suggest residual functioning of ToMM in some of these children (see Leslie and Roth, 1993; Happé, 1994; for discussion of this issue).

3.3.2.1. Correlations. In the ‘screen’ control task there is no calculation of an agent’s mental state. However, because one aspect of a situation must be selected in the face of intrusion by another salient and confusable aspect, the task may require SP. We do not pretend that the ‘screen’ task is an exact analogue of the problem solving structure of a standard false-belief task – its ‘parallelism’, arising from physical similarity, will be less potent than the parallelism of representation. We are not surprised to discover that screen is overall an easier task than standard false belief. Nevertheless, it seemed to produce a measurable stress upon the 3-year-old SP.

Our model states that normally developing children have their performance on standard false belief limited by SP functioning. This finds support in the highly significant correlation for the preschoolers between screen and standard false belief. On the other hand, SP should not limit performance on the non-standard false-belief task and accordingly we found no correlation between screen and search. For autistic children, our model extrapolates from their good performance on out-of-date photographs and maps (Leslie and Thaiss, 1992) to state that their selection processing is sufficiently powerful *not* to limit their performance on standard false belief. This predicts that there will be no distinction for autistic children between standard and non-standard false-belief tasks. As we saw earlier, this was confirmed by our finding that, in contrast to the normal case, non-standard false belief was not easier than standard. The model for autism receives further corroboration by our finding that, again in contrast to the normal case, screen did not correlate with standard false belief.

4. General discussion

Three-year-olds in experiment 1 did not find a true-belief task easier than they found its false-belief analog. The difficulty of belief tasks does not stem simply from the falseness of the attributed belief, but from the demands of correctly calculating an attributed belief content. In the ‘search’ task of experiment 2, we gave 3-year-olds ‘help’ with calculating a false-belief content and their performance improved considerably. The limiting factor in their performance on belief problems appears, at least in part, to be with a selection processing component shared with the non-theory of the mind screen task. However, we remind ourselves again of the caveat, expressed in footnote 2, that we cannot be sure at this stage of exactly how general selection processing is.

Previous research has demonstrated the well-known ‘theory of mind’ impairment in autistic children when compared with normally developing 4-year-olds. We were able to address a different question: given that children with autism and normally developing 3-year-olds both fail false-belief tasks, do they fail for the same reasons? In fact, the two groups show distinct patterns of performance. Autistic children are less likely to correctly attribute ‘knowledge’ but when they do, they are more likely to use that attribution to predict behavior. Three-year-olds by contrast attempt to predict behavior on the basis of an attributed belief rather than knowing/not knowing, even at the cost of failure. Autistic children also differ from 3-year-olds in not finding a non-standard false-belief task easier and in succeeding on control tasks, like photographs or the ‘screen’ task. The 3-year-old’s performance on belief tasks is limited by selection processing, the autistic child’s by an impaired ToMM.

As we pointed out earlier, the notion of task structure relates to the way a given task is processed by the cognitive system. Task structure then is not simply to be equated with experimental procedure. Rather it implicates a deeper and theoretical analysis of how a task is processed. Can we find other examples of false-belief tasks where a task analysis illuminates 3-year-old success? We think there are several examples in the literature. We briefly review four of them.

Roth and Leslie (1991) gave 3-year-olds a task in which Sally leaves an object out in the open before going away. Ann then hides the object in a box. When Sally returns, she immediately sees that the object is not where she left it. As a result, she asks Ann where it is. Ann replies deceptively, saying it is in a third location. The children were focused on the verbatim exchange between Sally and Ann which was acted out in direct speech and repeated for each child, checking each time that the child had attended to the conversation. The children were then asked about the beliefs of Sally and Ann with respect to the object’s location. It was found that the 3-year-olds tended to ascribe a false belief both to Sally and to Ann. Although the children were, strictly speaking, wrong to ascribe a false belief to Ann, the deceiver, nevertheless it was a false belief they ascribed. Autistic children, in contrast, tended to give reality-based answers when asked about both Sally's and Ann's beliefs. The relevant aspects of this task would seem to be (a) when Sally returns, she sees the object is not where she left it, (b) she immediately reacts to this by asking where it is, and (c) the other character in answering this question expresses the false content in behavior by (falsely) asserting another location. Although we are yet unable to specify the exact roles these features played individually or whether they interact, it seems plausible that the 3-year-olds were helped in inhibiting the inference from current reality by the first two features and in selecting the false content by the last.

Other tasks may help 3-year-olds by reducing the potency of the inference from reality. For example, in Wellman and Bartsch's 'not own' belief task, (Wellman and Bartsch, 1988) the child is told only of two *possible* locations, *x* and *y*, for the object and that a protagonist believes it is in *y*. Under these circumstances, 3-year-olds do well and predict belief-appropriate behavior for the protagonist. Similarly, 3-year-olds do well if they are only *told* about an object's location but do not get to see it for themselves (Zaitchik, 1991). In this experiment, an actor tells the child that the object is really in location *x* but then deceptively tells a second actor that the object is in location *y*. In these cases, the competing belief contents have more nearly equal potency, presumably because the child is not sure which belief is true, reducing the need for SP. Or consider Mitchell and Lacohée (1991). In that study, 3-year-olds were helped to recall their own false belief when first seeing a Smarties tube, which turned out later to contain a pencil, by 'mailing' a picture of Smarties prior to being shown the pencil. It appears that posting the picture of the ('falsely believed') Smarties helps to balance the potency of the true belief or to select the particular content of the false belief or both. In further studies Mitchell and others have reinforced and extended this analysis (see Freeman and Lacohée, 1995; Saltmarsh et al., 1995).

In summary, we believe that the 3-year-old can conceive of beliefs and their contents, even when false. But she has difficulty calculating accurate contents especially when the relative salience of a particular content is low compared with the salience of a content reflecting the current state of the world. This is normally not a problem because most often beliefs are true. But some situations defeat this simple assumption and with these unusual situations the 3-year-old has difficulty. Unless a belief problem inputs with the right kind of format, the younger child will not recognize it as a false belief. The main impact of this limitation is upon younger children's ability to learn about beliefs.

4.1. Concept acquisition in 'theory of mind'

Let us return to the larger questions concerning concept acquisition with which we began. First, we bring into focus a major difficulty in developing a 'theory-theory' of the acquisition of mental state concepts. We will then summarize some of the main ideas that guide our own approach.

The theory-theory holds that the possession or acquisition of ‘theory of mind’ concepts is a consequence of the possession or acquisition of a theory, specifically, of a theory of what mental states ‘really’ are. According to this view, it is the theory that gives the concept its meaning and, therefore, determines the identity of the concept. Consequently, the acquisition of the theory determines the acquisition of the concept. This general view of the nature of concepts and concept acquisition has been developed in domains other than ‘theory of mind’(e.g. Carey, 1985; Gopnik and Meltzoff, 1997). In the case of the concept, BELIEF, the prevailing assumption is that the concept is acquired by acquiring the ‘representational theory’. Perner has been at pains (e.g. Perner, 1991, 1995) to spell out what it is exactly the child must hold when she holds the ‘representational theory of mind’. He is right to have focused on this question, because spelling out the content of the child’s theory is the nub of the scientist’s theory of that theory. This is because, according to the theory-theory, in order to grasp the concept, BELIEF, it is necessary to grasp a *theory* of what a belief ‘really’ is[4].

Earlier we discussed Perner’s proposal for the content of the child’s representational theory of belief. Perner’s suggestion is that the (4-year-old) child acquires the concept, BELIEF, by arriving at the theory that *when John believes that* p, *John semantically evaluates a mental representation that expresses the proposition that* p. As we pointed out, stating the child’s theory inevitably means employing a number of other concepts (SEMANTIC EVALUATION, MENTAL REPRESENTATION, etc.). And because the theory is attributed to the child, these other concepts are of course attributed to the child too. But how does the child come to grasp *these* concepts, all of which are equally or more abstract than BELIEF? For a theory-theorist, the answer has to be either that these concepts are innate or that these concepts are grasped by grasping other theories. Then these further theories will spawn yet further sets of concepts which will either be innate or acquired by acquiring further theories. It is not clear where or how this will end. The problem would be less vicious if the iterations of attributed theories and concepts were constrained so that each set of concepts spawned had to be less abstract than the set before. This way one would hope to reduce eventually to a set of basic behavioral concepts, whose acquisition everyone at least felt they understood. However, because this reductionist avenue must cast the child’s ‘theory of mind’ as behavioristic, it is unappealing to anyone committed to accounting for its mentalistic character.

4.1.1. Learning about beliefs

The problems inherent in developing a theory-theory of BELIEF encourage us to explore different assumptions. We would like a theory of concept possession that can make sense of the fact that young children possess abstract concepts without entailing that they grasp arcane theories. In this vein, we have proposed that ToMM can be thought of as an attentional mechanism, specialized for selectively attending to mental states (Leslie and Roth, 1993; Leslie, 1994a). Consider the mechanisms of

[4]We focus our remarks on Perner because, although other theorists, like Wellman or Gopnik, also subscribe to the ‘representational’ theory-theory, they have not spelled out what they think is the content of the child’s representational theory of belief. Presumably they subscribe to a similar content.

color vision: don't these mechanisms allow children to attend to instances of *red* without children having to employ a theory of *what color is*? Is it possible that there might exist mechanisms which allow even abstract properties to be attended to without employing theories of what these properties 'really' are? We suggest there are. Specifically, ToMM allows children to recognize pretense in others without children employing a theory of *what pretending is*. Leslie and Happé (1989) suggested one way in which this might happen. They speculated that ToMM targets the PRETEND concept to the characteristic demeanor of someone who is pretend playing: the exaggerated gestures and intonation, the 'sound effects', the 'knowing' looks and smiles, and so forth, that have been noted by observers of pretense (e.g. Piaget, 1962; Bateson, 1972; Huttenlocher and Higgins, 1978; Leslie, 1987). Although these signals cannot constitute the meaning of the concept, PRETEND, which ToMM employs, watching an agent make such a display may trigger ToMM to hypothesize that the agent is pretending and to infer a content for the pretense. Given the intentionally communicative nature of such signals on the part of the pretender, targeting pretend M-representations in this way may have a high degree of success. Asking about target signals for a concept can be aligned with information-based approaches to concepts (e.g. Fodor, 1997; Margolis, in press).

The emergence of the M-representation solves a critical but limited problem. It allows the young brain to *attend* to the mental states of agents despite the fact that such states cannot be seen, heard, felt, or otherwise sensed. Leslie (1994a) described this as 'the fundamental problem of theory of mind' because being able to attend to mental states is a prerequisite for learning about those states. ToMM solves this problem using the M-representation: ToMM attends to agents' behavior and tries to infer the mental state from which the behavior issues and which gives the behavior its meaning. To avoid misunderstanding: though we claim ToMM allows the child to attend to (some) mental states, we are *not* thereby claiming that ToMM has or uses the concept, MENTAL STATE, to do this. Belief-metarepresentations give the child a critical but limited power – they allow the child to attend to beliefs and therefore to *learn* about them. Though ToMM introduces concepts like PRETEND and BELIEVE, ToMM does not have, or at any rate does not need, any particular knowledge about pretending or believing. Specifically, the child does not need to possess or develop a theory of *what a belief really is* in order to recognize belief in others. But, with the ability to attend to pretending or to believing, knowledge about these states can be acquired, by ToMM or by other systems.

The well-known developmental pattern of error on false-belief tasks reflects, we suggest, a general characteristic of the immature cognitive system. The error is to fail to inhibit a default and normally correct assumption, in this case, that the content of a belief is true. This error then makes it difficult for the young child to attend to false beliefs. This in turn makes it difficult for the child to learn things about false beliefs. However, in *some* circumstances, even the young child can succeed in attending to a belief-with-a-false-content. This limited but critical ability leaves open the door to learning. With maturation, learning, and practice the child becomes more able to inhibit the 'truth assumption'. Eventually, attention will readily target false beliefs; the false belief in a scenario will 'pop out'. With conceptual pop-out, the child can rapidly learn about beliefs, e.g. that beliefs tend to be false under systematic circumstances, that surprise occurs when a false belief is relieved, that false beliefs can be induced in opponents to useful effect, how to do the latter and get away with it, and so on. By contrast, in autistic disorder, an impaired ToMM means that it will remain hard to attend to and therefore to learn things about belief and other mental states, impairing the autistic child's social skills.

Any theory of the possession and acquisition of abstract concepts faces great difficulties. But our account seems simple in comparison with 'theory-theory' alternatives. We do not expect a young child to discover a theory of what beliefs 'really' are in order to recognize a false belief. Instead, the brain includes a mechanism which bestows a specific conceptual competence, namely, the concept, BELIEF, plus a set of processes that leverage this initial competence. Together, these resources allow the child to accumulate knowledge about beliefs.

Acknowledgements

We are grateful to Pam Polizzi, Tim German, Stephen Stich and Jerry Fodor for helpful discussions and comments on earlier drafts and to two anonymous reviewers. We thank Laila Thaiss for help in collecting parts of the data in experiment 1. We also thank Christchurch School, Dorset, England, Sybil Elgar School and Whitefields School, London, and Tel Aviv nursery schools for their help.

References

Astington, J., Harris, P., Olson, D., 1988. Developing Theories of Mind. Cambridge University Press, Cambridge.

Baron-Cohen, S., 1990. Autism: a specific cognitive disorder of 'mind-blindness'. International Review of Psychiatry 2, 79–88.

Baron-Cohen, S. (1995). Mindblindness: An Essay on Autism and Theory of Mind. MIT Press, Cambridge, MA.

Baron-Cohen, S., Leslie, A.M., Frith, U., 1985. Does the autistic child have a 'theory of mind'? Cognition 21, 37–46.

Baron-Cohen, S., Tager-Flusberg, H., Cohen, D., 1993. Understanding Other Minds: Perspectives From Autism. Oxford University Press, Oxford.

Bateson, G., 1972. A theory of play and fantasy. In: Bateson, G. (Ed.), Steps to an Ecology of Mind. Chandler, Aylesbury, pp. 177–193.

Butterworth, G., 1993. Towards an ecology of mind. Commentary. Behavioral and Brain Sciences, 16, 31–32.

Butterworth, G., Harris, P.L., Leslie, A.M., Wellman, H.M., 1991. Perspectives on the Child's Theory of Mind. Oxford University Press, Oxford.

Carey, S., 1985. Conceptual Change in Childhood. MIT Press, Cambridge, MA.

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

Charman, T., Baron-Cohen, S., 1992. Understanding drawings and beliefs: a further test of the metarepresentation theory of autism (Research Note). Journal of Child Psychology and Psychiatry 33, 1105–1112.

Charman, T., Baron-Cohen, S., 1995. Understanding photos, models, and beliefs: a test of the modularity thesis of theory of mind. Cognitive Development 10, 287–298.

Chomsky, N.A., 1965. Aspects of the Theory of Syntax. MIT Press, Cambridge, MA.

Diamond, A. 1988. Differences between adult and infant cognition: is the crucial variable presence or absence of language? In: Weiskrantz, L. (Ed.), Thought Without Language. Oxford Science Publications, Oxford, pp. 335–370.

Flavell, J.H., 1988. The development of children's knowledge about the mind: from cognitive connections to mental representations. In: Astington, J.W., Harris, P.L., Olson, D.R. (Eds.), Developing Theories of Mind. Cambridge University Press, New York, pp. 244–267.

Fodor, J.A., 1976. The Language of Thought. Harvester Press, Hassocks, Sussex.

Fodor, J.A., 1981. Propositional attitudes. In: Fodor, J.A. (Ed.), Representations: Philosophical Essays on the Foundations of Cognitive Science. Harvester Press, Brighton, pp. 177–203.

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

Fodor, J.A., 1997. Concepts: Where Cognitive Science Went Wrong. Oxford University Press, Oxford.

Freeman, N.H., Lacohée, H., 1995. Making explicit 3-year-olds' implicit competence with their own false beliefs. Cognition 56, 31–60.

Frith, U., 1989. Autism: Explaining the Enigma. Blackwell, Oxford.

Frith, U., Morton, J., Leslie, A.M., 1991. The cognitive basis of a biological disorder: autism. Trends in Neurosciences 14, 433–438.

Gopnik, A., 1993. How we know our minds: the illusion of first-person knowledge of intentionality. Behavioral and Brain Sciences 16, 1–14.

Gopnik, A., Meltzoff, A.N., 1997. Words, thoughts, and theories. MIT Press, Cambridge, MA.

Gopnik, A., Wellman, H.M., 1995. Why the child's theory of mind really is a theory. In: Davies, M., Stone, T. (Eds.), Folk Psychology: The Theory of Mind Debate. Blackwell, Oxford, pp. 232–258.

Happé, F.G., 1994. An advanced test of theory of mind: understanding of story characters' thoughts and feelings by able autistic, mentally handicapped, and normal children and adults. Journal of Autism and Developmental Disorders 24, 129–154.

Happé, F.G., 1995. Autism: an introduction to psychological theory. Harvard University Press, Cambridge, MA.

Hogrefe, G.J., Wimmer, H., Perner, J., 1986. Ignorance versus false belief: a developmental lag in attribution of epistemic states. Child Development 57, 567–582.

Huttenlocher, J., Higgins, E.T., 1978. Issues in the study of symbolic development. In: Collins, W. (Ed.), Minnesota Symposia on Child Psychology, Vol. 11. Erlbaum, Hillsdale, NJ, pp. 98–140.

Julesz, B., 1991. Early vision and focal attention. Reviews of Modern Physics 63, 735–772.

Julesz, B., 1994. Dialogues on Perception. MIT Press, Cambridge, MA.

Leekam, S., Perner, J., 1991. Does the autistic child have a 'metarepresentational' deficit?. Cognition 40, 203–218.

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

Leslie, A.M., 1988a. Some implications of pretense for mechanisms underlying the child's theory of mind. In: Astington, J., Harris, P., Olson, D. (Eds.), Developing Theories of Mind. Cambridge University Press, Cambridge, pp. 19–46.

Leslie, A.M., 1988b. The necessity of illusion: perception and thought in infancy. In: Weiskrantz, L. (Ed.), Thought Without Language. Clarendon Press/Oxford University Press, Oxford, pp. 185–210.

Leslie, A.M., 1991. The theory of mind impairment in autism: evidence for a modular mechanism of development? In: Whiten, A. (Ed.), Natural Theories of Mind: Evolution, Development and Simulation of Everyday Mindreading. Blackwell, Oxford, pp. 63–78.

Leslie, A.M., 1992. Autism and the 'Theory of Mind' module. Current Directions in Psychological Science 1, 18–21.

Leslie, A.M., 1994a. Pretending and believing: issues in the theory of ToMM. Cognition 50, 211–238. Reprinted in: Mehler, J., Franck, S. (Eds.), Cognition on Cognition, 1995. MIT Press, Cambridge, MA, pp. 193–220.

Leslie, A.M., 1994b. ToMM, ToBy, and Agency: Core architecture and domain specificity. In: Hirschfeld, L., Gelman, S. (Eds.), Mapping the Mind: Domain Specificity in Cognition and Culture. Cambridge University Press, New York, pp. 119–148.

Leslie, A.M., Frith, U., 1988. Autistic children's understanding of seeing, knowing and believing. British Journal of Developmental Psychology 6, 315–324.

Leslie, A.M., Frith, U., 1990. Prospects for a cognitive neuropsychology of autism: Hobson's choice. Psychological Review 97, 122–131.

Leslie, A.M., German, T.P., 1995. Knowledge and ability in 'theory of mind': one-eyed overview of a debate. In: Davies, M., Stone, T. (Eds.), Mental Simulation: Philosophical and Psychological Essays. Blackwell, Oxford, pp. 123–150.

Leslie, A.M., German, T.P., Happé, F.G., 1993. Even a theory-theory needs information processing: ToMM, an alternative theory-theory of the child's theory of mind. Commentary. Behavioral and Brain Sciences 16, 56–57.

Leslie, A.M., Happé, F., 1989. Autism and ostensive communication: the relevance of metarepresentation. Development and Psychopathology 1, 205–212.

Leslie, A.M., Polizzi, P., in press. Inhibitory processes in the false belief task: two conjectures. Developmental Science.

Leslie, A.M., Roth, D., 1993. What autism teaches us about metarepresentation. In: Baron-Cohen, S., Tager-Flusberg, H. Cohen, D. (Eds.), Understanding Other Minds: Perspectives from Autism, Oxford University Press, Oxford, pp. 83–111.

Leslie, A.M., Thaiss, L., 1992. Domain specificity in conceptual development: neuropsychological evidence from autism. Cognition 43, 225–251.

Margolis, E., in press. How to acquire a concept. Mind and Language.

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

Perner, J., 1988. Developing semantics for theories of mind: from propositional attitudes to mental representation. In: Astington, J., Harris, P.L., Olson, D. (Eds.), Developing Theories of Mind. Cambridge University Press, Cambridge, pp. 141–172.

Perner, J., 1989. Is 'thinking' belief? Reply to Wellman and Bartsch. Cognition 33, 315–319.

Perner, J., 1991. Understanding the Representational Mind. MIT Press, Cambridge, MA.

Perner, J., 1995. The many faces of belief: reflections on Fodor's and the child's theory of mind. Cognition 57, 241–269.

Perner, J., Baker, S., Hutton, D., 1994. Prelief: the conceptual origins of belief and pretense. In: Lewis, C., Mitchell, P. (Eds.), Children's Early Understanding of Mind: Origins and Development. Lawrence Erlbaum, Hillsdale, NJ.

Piaget, J., 1962. Play, Dreams and Imitation in Childhood. Routledge and Kegan Paul, London.

Pinker, S., 1994. The Language Instinct: How the Mind Creates Language. Morrow, New York.

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.

Roth, D., 1993. Beliefs about false beliefs: understanding mental states in normal and abnormal development. Ph.D. Thesis, Tel Aviv University.

Roth, D., Leslie, A.M., 1991. The recognition of attitude conveyed by utterance: a study of preschool and autistic children. British Journal of Developmental Psychology 9, 315–330. Reprinted in: Butterworth, G.E., Harris, P.L., Leslie, A.M. Wellman, H.M. (Eds.), Perspectives on the Child's Theory of Mind. Oxford University Press, Oxford, pp. 315–330.

Russell, J., Mauthner, N., Sharpe, S., Tidswell, T., 1991. The 'windows task' as a measure of strategic deception in preschoolers and autistic subjects. British Journal of Developmental Psychology 9, 331–349. Reprinted in: Butterworth, G.E., Harris, P.L., Leslie, A.M., Wellman, H.M. (Eds.), Perspectives on the Child's Theory of Mind. Oxford University Press, Oxford, pp. 331–349.

Saltmarsh, R., Mitchell, P., Robinson, E., 1995. Realism and children's early grasp of mental representation: belief-based judgements in the state change task. Cognition 57, 297–325.

Sullivan, K., Winner, E., 1991. When 3-year-olds understand ignorance, false belief and representational change. British Journal of Developmental Psychology 9, 159–172. Reprinted in: Butterworth, G.E., Harris, P.L., Leslie, A.M., Wellman, H.M. (Eds.), Perspectives on the Child's Theory of Mind. Oxford University Press, Oxford, pp. 331–349.

Treisman, A., 1988. Features and objects: the 14th Bartlett memorial lecture. Quarterly Journal of Experimental Psychology 40, 201–237.

Treisman, A., Gelade, G., 1980. A feature integration theory of attention. Cognitive Psychology 12, 97–136.

Wellman, H.M., 1990. The Child's Theory of Mind. MIT Press, Cambridge, MA.

Wellman, H.M., Bartsch, K., 1988. Young children's reasoning about beliefs. Cognition 30, 239–277.

Whiten, A., 1991. Natural Theories of Mind: Evolution, Development and Simulation of Everyday Mind-reading. Basil Blackwell, Oxford.

Wimmer, H., Hogrefe, G.J., Sodian, B., 1988. A second stage in children's conception of mental life: understanding sources of information. In: Astington, J., Harris, P.L., Olson, D.R. (Eds.), Developing Theories of Mind. Cambridge University Press, New York.

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.

Zaitchik, D., 1990. When representations conflict with reality: the preschooler's problem with false beliefs and 'false' photographs. Cognition 35, 41–68.

Zaitchik, D., 1991. Is only seeing really believing?: sources of true belief in the false belief task. Cognitive Development 6, 91–103.