Mental Files in Development: Dual Naming, False Belief, Identity and Intensionality

Josef Perner · Brian Leahy

Abstract We use mental files to present an analysis of children's developing understanding of identity in alternative naming tasks and belief. The core assumption is that younger children below the age of about 4 years create different files for an object depending on how the object is individuated (e.g., as a rabbit or as an animal). They can anchor them to the same object, hence think of the same object whether they think of it as a rabbit or as an animal. However, the claim is, they cannot yet link their files to one another to represent that they have the same referent. Without linking the information contained in one file is not available in the other file. Hence, when thinking of the object as a rabbit (using the rabbit file) the information that it is also an animal is not available. For representing a person's belief about an object a vicarious file contains what the person believes about the object. To capture that the belief is about that object the vicarious file has to be linked to the regular file, which by assumption children younger than 4 years cannot do. This assumption can therefore explain why problems with alternative naming and understanding false beliefs are overcome at the same age.

Our prime objective is to establish mental files as a useful tool for cognitive development. We think that they are useful because they capture important aspects of cognition that, at least in psychology, have been largely ignored, and because in our own research mental files proved helpful in explaining children's curious responses on developmental tests which they tend to master at about the age of 4 years. The phenomena pertain to tasks that involve differences of perspective, multiple labels and identities of objects, false beliefs, intensionality, episodic remembering—to name a few.

A useful first step in our programme lays out the new and useful features that mental files bring to cognitive analysis. Then we show why these features are helpful in illuminating the developmental phenomena we are in the short term interested in. To show that mental files can indeed fulfil some of these promises we go into depth of how

J. Perner (✉)
Centre for Cognitive Neuroscience, University of Salzburg, Salzburg, Austria
e-mail: josef.perner@sbg.ac.at

B. Leahy
Department of Linguistics, University of Konstanz, Constance, Germany

they can make us understand why children have problems with alternative naming, identity statements and why these problems relate to their problems understanding false beliefs.

Our point of departure is a body of earlier studies that addressed the same problem (e.g., (Perner and Brandl 2005), (Perner et al. 2007)) and which were influenced in their theorizing by work on discourse referents by Hans Kamp (1990) and Irene Heim's (2002) file management theory. This paper takes a narrower empirical focus. It provides more details about how the theory explains the data that is most pivotal for us: the relation between success on false belief and alternative naming tasks. It also brings the analysis in line with the role of mental files in thinking more generally (Recanati 2012).

In the next section (1) we describe some of the critical data, the alternative naming (AN) data. In section (2) we describe some critical features of mental files. In section (3) we describe mental files at the interface between discourse and perception; section (4) expands on what we mean with talk of identifying referents. Section (5) uses the tools developed so far to explain why young children fail the alternative naming task, while (6) describes the difference between younger and older children such that older children can pass the alternative naming task. Section (7) uses mental files to explain the difference between children who fail and children who pass the false belief task. (8) overviews the developmental theory we see forming and (9) concludes.

1 The Alternative Naming Game

The alternative naming (AN) data are generated as follows. Test subjects are first administered a vocabulary test that establishes that they know multiple kind-terms for a single object that are either synonyms (e.g., rabbit-bunny) or that stand in kind-subkind relations (e.g., rabbit-animal). For example, they were shown a piece of paper on which several objects were drawn but only one was a rabbit; the rabbit was moreover the only animal. Subjects were asked to point out the rabbit; later they were shown the same array and were asked to point out the bunny or the animal. It was inferred that children who succeed in these tasks know that the object is a rabbit and know that it is a bunny or know that it is an animal.

Test subjects are then presented with the alternative naming task. In one version, the experimenter uses the same set of drawings and says something along the lines of (Doherty 1994, Experiment 1):

"This is Tony, and he is going to tell us what these things are. But Puppet also wants to play but he doesn't want to say the same thing as Tony, so we are going to help him think of a different way of saying what these things are. Now, Tony says this is a rabbit. What could puppet say it is? I know, he could say it is a bunny, because it is a rabbit and it is a bunny. They are different ways of saying the same thing.

Test: Tony says: "this is a rabbit." What could Puppet say it is? < child's answer>

If the experimenter says "Rabbit", then, a child who says "Bunny" or "Animal" plays correctly. Three-year-olds typically fail this task; they most often repeat the word

produced by the experimenter. Four-year-olds typically pass. Moreover, passing this test is strongly correlated with passing the false belief task.

The AN task is supplemented with various controls. On the two colour control (Perner et al. 2002), the experimenter uses a set of drawings of objects with two colours, among them a flower with red petals and a green stem. The experimenter says something along the lines of, "Here are some things with two colours. We can play a game. It goes like this: I say a colour, and you have to say the other colour of the same thing, not the one that I said. Look: the flower is red and green. Now I say red, and you have to say the other colour, not the one that I just said." The experimenter then offers one of the two colours; then comes the test question: Now you have to say the other colour!

If the experimenter says "Red", then, a child who says "Green" plays correctly. Three- and four-year-olds find this task easy, whether or not they pass the false belief task.

The first questions we will ask are, "Why are three-year-olds successful in the two colour control and not in the alternative naming task? Why is the two colour task easier than the alternative naming task?". In any case, the fact that the children who have problems with the alternative naming task have no serious problems with the two colour control task helps rule out obvious alternative interpretations of their difficulty with the alternative naming task. In particular, their difficulty can't be a problem of inhibiting repetition of what the partner has said.[1] Then we need to expand that theory to explain why success on the alternative naming task so strongly coincides with success on the false belief task. We will discuss some additional useful tasks along the way.

In the alternative naming task, the experimenter says, "I call it a rabbit". In the two colour control, the experimenter says, "Now I say it is red". The child is then charged to "say the other word, not the one that I just said" or "say the other colour, not the one that I just said". What is the relevant difference between these? In the two colour control, the experimenter has stated a property of the object ostended, and the test subject must state another property of that object. In the alternative naming task, the experimenter has used an individuating term that helps the subject determine as-what an ostended individual is being talked about, and the child must state a second term that individuates the same object in a different way, that is, that individuates the same object as something else. So children who fail the test have a hard time stating a second kind to which an object belongs, though they can easily state a second property that the object displays.

Note that children who fail the task do not have a problem stating a first kind to which an object belongs. Asked simply to say what an ostended rabbit is, children who fail the alternative naming task but have adequate vocabulary typically reply 'bunny' or 'rabbit'. The problem arises when they are asked to state a second kind to which the object belongs, when the discourse has already established a kind to which the object

---

[1] Or as one of the reviewers has pointed out that the repeated emphasis on not repeating what the experimenter has said may lead to a pink elephant phenomenon that the children trying to avoid repetition can't help say what they are intent on avoiding.

belongs.[2] And properties do not seem to raise the same problem. We now sketch our explanation for why this should be. It turns on an empirical hypothesis regarding how information is stored in the brain. It will make much of the difference between kind terms and terms for properties, claiming that information provided by these terms is stored in different ways when they are employed in the manner observed in the alternative naming task. We now engage in a substantial discussion to spell out the distinction between kind terms and property terms, which we see in the role they play in identifying the external object talked about, i.e., the referent.

2 Critical Features of Mental Files

At their most basic, mental files are devices to *track* something (say, its *target*) and to *store information* about it. This is very basic and is a necessary feature of many things, e.g.,: concepts, predication, discourse coherence, recognition, following an object, and so on. For instance, for Millikan (2000) to have a concept of something is to have a file for it, to be able to track it and accumulate information about it so as to be able to reidentify the object at later times and apply knowledge gained through previous encounters. We can have concepts of individuals (my cat, Santa Claus), of substances (water, gold, phlogiston), of properties (red, soft, round), etc. Files capture the predicative structure of language, i.e., the distinction between what one is talking about (the subject, topic, i.e., what the file tracks) and what one says about it (the information about the topic, i.e., the information the file has on it). This is also an important feature of our thinking (at least language like cognitions, language of thought): what we are thinking of and what we know about it.

Files are critical for discourse coherence. Repeated mention of the same entity needs to be mapped onto the same mental file so that information given about this entity accumulates and is not lost in unrelated places of storage. Karttunen (1976) introduced the notion of discourse referent for that purpose and to move from the logical analysis of sentences to the logical structure of discourse (Kamp and Reyle 1993). Language uses its specific devices for making clear which entity is meant when a referring expression occurs. Typically use of the same expression (same proper or common name) indicates the same individual for which the name had been used before but referring with pronouns requires more subtle rules of anaphoric reference. An analogue problem exists with visual perception. Seeing the same entity again requires means of recognizing it as the same as seen earlier. Again same appearance (like same name) is a basic guide to identity, but given that moving and flexible entities change their appearance drastically between encounters visual scientists face a nontrivial task specifying our mental processes that allow for recognition under such difficult circumstances. Beside visual recognition facilities a mental file also needs to be able to track

---

[2] One may well wonder what children, who despite the very explicit instructions not to repeat partner's answer still repeat that very label, do understand of the instructions. We don't really know. They probably try to be cooperative and make the best of something they don't fully comprehend. They most likely pick up on "having to say what the ostended object is," and then cannot but repeat what the experimenter had just called it. Interestingly, when instructed to pretend the object was something else and, therefore, call the object something it is not (control condition in Experiment 4 of Doherty and Perner 1998) children had few problems.

its object when it is moving without having to re-recognize it again every split second. Pylyshyn (2007) has proposed the existence of fingers of instantiation (FINSTs) for that purpose.

Our prime objective is to make use of mental files to explain developmental data about alternative naming, understanding identity statements, dual identity, and false beliefs. For this we focus on the interplay between language and perception.

3 Mental Files in the Interface of Language and Perception

When talking about an object in the perceptual field one has to first of all individuate the object—determine what one is talking about. In order to cut a complicated topic short we will concentrate on the type of situations relevant for our developmental work on alternative naming and illustrate some of the critical distinctions. For instance, when pointing at a dancing bear and saying "this is a lovely bear" the listener creates a file for the bear and encodes on it that it is a bear and that it is lovely, and perhaps some other evident features. The file has to be anchored (Kamp 1990) to the object, so that observed changes in the object can be registered in the appropriate file. This can be achieved in the immediate context by something like Pylyshyn's (2007) FINSTs, direct links between an object file and it external referent that allow tracking of moving objects. For more long term tracking the file needs to have detailed perceptual information to be able to recognize the individual object and discriminate it from other objects.

An interesting contrasting case to using "bear" occurs when using the same pointing gesture (pointing toward the bear) but one says "This is a lovely tango." In this case one picks out the bear's movements and not the bear and hence a different file is to be created.

Another critical contrast case for us arises when pointing at the bear and saying "This is a lovely bear" or "lovely animal". In both cases the same physical object is picked out and on those grounds one theoretical option (same file) is that the listener should create essentially the same file in both cases.

Another theoretical option (different files) is that the listener creates different files in the two cases. This option tries to capture two intuitions, namely that if an object is individuated by a different name it gets a different "identity" (Markman 1989[3]; Flavell 1988), that by giving an object a different label one puts a different perspective on it (Clark 1997; Tomasello 1999), and that one can think of the same object in different ways, e.g., as a bear or as an animal. To capture these intuitions we assume that the

---

[3] Markman (1989) expressed this view in the context of explaining mutual exclusivity phenomena: "children may believe that an object has one and only one identity – that it can only be one kind of thing – and that an object's identity is revealed by object labels" (Markman 1989, p. 212); Similarly John Flavell: "the child may think that the word *kitty* describes or characterizes the single way that cats are, and consequently that calling them "cat" or "animal" amounts to claiming that they are some other way than that. It is possible, therefore, that young children are loath to accept two category names for the same thing, for much the same reason they are loath to accept two perspectives, or both the appearance and the reality, as two equally valid character-izations of the same thing: namely, because one thing has only one nature and should therefore be characterized in only one way." (Flavell 1988, p. 254).

listener creates different files in the two cases and gives them a label or heading "bear" or "animal", respectively.

Understanding that thinking of the bear as a bear and thinking of it as an animal means that one is thinking about the same entity can be achieved by two measures. The two files need to be anchored to the same external object and the fact that they are anchored to the same object needs to be represented. How this can be and should be built into the theory we will address later on. For now we have reached a first point where we can state an interesting developmental claim: children younger than about 4 years can anchor each file to the same object but cannot represent this fact, i.e., cannot represent the identity of the object referred to by the files. In other words, although these children represent the same object with different files, they have no awareness of that. From their perspective the bear as bear and the bear as animal are two different objects, just like pointing to the bear and saying "This is a bear" or saying "This is a tango" creates different files for different objects of which different properties may be predicated. This theoretical proposal provides an explanation of young children's curious inability to play the alternative naming game.

Because this point is of central importance, let us try to explain it in a different way. We maintain that individuating an object requires individuating it *as* something. For example, when someone ostends an object and says, "This animal has soft fur", she individuates the object as an animal. When she ostends the same object and says, "This rabbit has long ears", she individuates the same object, the same individual. But she individuates it *as* something else, under a different perspective. When a single object can be individuated as multiple different things, the way it is being individuated in a conversation needs to be tracked for at least as long as the object is being individuated in that way. If an error is made in this respect, that is, if a speaker inadequately tracks as-what an object is being individuated in a conversation, there is a risk of confusion arising. If one conversant individuates an object as an animal and another individuates it as a rabbit, they risk misinterpreting some of each other's utterances.

For this reason we propose that a mental file for an object under discussion always contains information about as-what that object is being individuated. An object that is under discussion can only be individuated *as* one thing at any given moment in a conversation; consequently, a single mental file, at any given time, can only reflect one way that object is being individuated. So if a person can individuate (not in conversation, but privately, mentally) a single object as more than one thing at a single moment in time, then that person must have two or more mental files for that object, each incorporating information about that object *as* something different (for example, one file incorporating information about that object as a rabbit, another, as an animal). We will say that the files (and with them their respective object) must be *labelled* differently.

We have claimed that individuation requires individuation-as. We provided some examples where a speaker used a linguistic term ('rabbit' or 'animal') to make clear as-what an ostended object should be individuated. But when we make claims about objects around us, we do not always explicitly provide linguistic terms that make clear as-what the subject of our discourse should be individuated. When we do not use any such explicit linguistic term, we trust that contextual information is sufficiently rich for our interlocutors to determine as-what we are individuating the objects we individuate.

A mental file for an observed object always incorporates information about as-what an object under discussion is being individuated.

This last fact is important because it demonstrates how a mental file always encodes information about its referent from a particular perspective. Even if language has not explicitly fixed a perspective, some perspective is fixed. This perspective-relativity of files is essential to the explanatory power of our theory. First, it helps us explain a difference between the behaviour of noun phrases and adjectives observed in the alternative naming task. We turn to discussing that behaviour now. Second, it will contribute to our explanation of the data gathered from the alternative naming and false belief tasks, discussed thereafter.

4 Identifying the Referent

We need to show two things: (1) the use of noun phrases (individuating terms) in the alternative naming task plays a different role in individuating the referent than the use of adjectives; (2) use of different individuating noun phrases leads to different mental files (alternative filing) but use of different adjectives does not.

Role in Individuation By just pointing at something and saying "This!" one has to rely on the listener's intuition of what one might mean. As we mentioned above, information about as-what the ostended object is being individuated needs to be provided pragmatically, since no overt verbal behaviour fixes as-what the ostended object is to be individuated. Pointing together with a noun phrase, e.g., "This is **a rabbit**," makes clear as-what the ostended object is to be individuated, provided there is only one object within the range of the point, that is, provided the point sufficiently determines an object to be individuated. For lack of a better expression we call the nouns that serve this function "individuating expressions" because of their important role in determining as-what the referent is individuated.

Nouns can also play a role that adjectives cannot play so well in determining what is being individuated. Using an adjective denoting the property of the intended referent, like "this is white," or "lovely", may help narrow the range of intuitions but will not give a decisive answer on its own with respect to as-what the object is being individuated or, in some cases, in determining what is being individuated. For even when there is only one white or lovely rabbit to be seen, the referent could be the rabbit, or the white fur of the rabbit, or some spot on the fur, etc.[4] So, the use of an adjective suggests that one is attributing a property to some identified or to be identified referent. The sentence 'This is white' does not explicitly tell us as-what the extended object is to be individuated; 'This is a rabbit' does.

Double Filing A critical feature of the alternative naming game is that the same referent is named repeatedly and differently, i.e., individuated in different ways, as different things: "This is a rabbit" and "This is an animal," or respectively "We called this a

---

[4] In the case of "white" another interpretation is possible, namely that one is not talking about the rabbit but about its colour, i.e., "This is white" means "the colour I am pointing at is white." Although possible it does not seem to occur in the contexts used in the alternative naming game.

rabbit," and "We called this an animal." In order to determine the referent the listener has to apply some identification criteria for each term. In the case of "rabbit" she may look for characteristic rabbit features, in the case of "animal" for more general features shared by all animals. These features may, and probably will overlap, but there is no necessity in the identification procedure to enforce the idea that one must be talking about the same thing. Hence it would be a wise precaution (by an automatic and preconceptual identification procedure) to provide different files even though both happen to refer to the same external entity.[5] The realization that "rabbit" and "animal" can—and in this instance, must—refer to the same entity could be left to a later reasoning step. For instance, since all rabbits are animals, a rabbit and an animal in the same location must be the same entity, or this rabbit could also be individuated as an animal, a pet, etc.

In the case of the two colour control condition where two colours of the same entity are named, there is no similar reason to create two different files. The statement "This is white," as we explained above, does not determine the referent in as direct a way as using an individuating term. It helps limit ambiguity. If there happens to be a white rabbit and a black cat within the pointing range then mentioning "white" helps make clear that the more likely referent is the rabbit and not the cat. It still leaves open whether one talks about the rabbit, the fur, a spot on the fur, etc. And this is even clearer when the formulation in the actual experiment is used, e.g., using the example with the red rose on a green stem: "Here are some things with two colours. We can play a game. It goes like this: I say a colour, and you have to say the other colour of the same thing, not the one that I said." So when the experimenter, pointing to the rose, says "This is red" and the child is to name the other colour, green, it is clear that the referent is the rose and it is red and also green. There is no inducement to create a file for a green flower and another for a red flower.

With these clarifications in place we can now give a possible explanation for the developmental course of children's ability to play the alternative naming game. We have to first explain why younger children find this game difficult and then what older children acquire that alleviates these problems. To encode the fact that every mental file is associated with exactly one perspective, and reflecting the fact that the perspective of a file can be fixed by an overt individuating expression, we will say that every file is *labelled* with an *individuating expression*, which is the expression that was used in individuating that object, if one was provided linguistically, or an expression that is pragmatically inferred, if none was linguistically provided.

5 Mental Files and Alternative Naming: Why Children Fail

The pretest of the alternative naming experiments shows that children can individuate the objects used in both relevant ways, e.g., as a *rabbit* and as an *animal*. However, in the test when Tony says "This is a rabbit," the object is individuated as a rabbit. This, by assumption, leads to a file labelled "rabbit" anchored to the external object. Now in

---

[5] This is not to be understood as an argument that it has to be that way, but that alternative filing is a possible and perhaps sensible option.

the test question "What could puppet say it is?" the "it" refers to what has been talked about, namely the entity which has just been individuated as a rabbit. Now, for the child who does not represent the fact that multiple files are anchored to the same object, whose files are not linked to one another, alternative files amount to representing different objects (that happen to spatio-temporally overlap), just like in the example with the dancing bear, where saying "This is a bear" and "This is a tango" leads to different files representing different referents. Hence, for the child, file "rabbit" and file "animal" represent different entities, which are identical but the children don't yet realize that. So when asked what else the object is, the child uses the 'rabbit' file to think about this object, since it has just been individuated as a rabbit, i.e., since the "rabbit" file is activated. The object already individuated as a rabbit cannot be anything but a rabbit. So children answer with "rabbit," counter to the instructions given. This explains their difficulty and wrong answers in the alternative naming game.

This explanation was framed for the instructions used originally by Doherty (1994) Experiment 1. Most of the later versions, however, used instructions in terms of object naming (e.g., Stummer 2001; see Perner et al. 2002):

> We can play a game. It goes like this: I say a name, and you have to say the other name, not the one that I said. Look: this can be either a rabbit or an animal. I call it _____ [the blank is filled in with an experimental item, for example with 'rabbit'], and you have to say the other word, not the one that I just said." The experimenter then offers one of the two words; then comes the test question: Now you have to say the other name!

Our explanation also applies here under the plausible assumption that (at least for children) to name an object is to individuate it in a certain way. Therefore an object just labelled "rabbit" cannot be labeled (correctly) in any other way.

6 Mental Files and Alternative Naming: Why Children Pass

Our proposal is that children around the age of 4 years link files that are anchored to the same object to one another. *Linking* is an operation of connecting different files so that the information contained in each file can flow freely to any other file linked to it (Recanati 2012, p 14, following Perry's 2002 terminology). [6] With this setup in place children can follow instructions in the alternative naming game.

Starting with the modelling phase, children are told: "… Now, Tony says this is a rabbit. What could puppet say it is? I know, he could say it is an animal, because it is a rabbit and it is an animal. They are different ways of saying the same thing." With these instructions, children build two files for the object, one labelled "rabbit" and the other "animal". Each file is anchored to the ostended object. The two files are then linked to enable flow of information between them. In the test phase, when Tony calls the object,

---

[6] Linking of files is not restricted to only two files. Our account is, therefore, not limited to two alternative labels in the alternative naming game. In fact, children could have individuated the objects in other ways, e.g., as a *pet*, a *living thing*, etc. This would not impinge on our theoretical claims. It just turns out that children this age typically do not venture beyond the terms that have been used in the pretest.

e.g., an animal, children activate[7] their discourse file "animal" for the object. When then asked: "What could Puppet say it is?" they can implement the earlier instructions that puppet should not repeat what Tony had said. Free flow of information among linked files enables children to access what else the *rabbit* is by looking for information about what it is in the linked files. In the file headed "animal" they will find that it is also an animal and can use this fact to answer the question.

7 Mental Files and False Belief

In the standard false belief task (Wimmer and Perner 1983) children are told a story in which a protagonist, e.g., Mistaken Max puts a chocolate bar into the kitchen table drawer (location 1) and leaves the scene. In his absence his mother unexpectedly moves the chocolate to the cupboard (location 2). When Max returns to get his chocolate, children are asked where he will go to get it, the drawer or the cupboard. Most children younger than 4 years answer with location 2, the actual location of the chocolate, while most children older than four answer correctly with location 1, where Mistaken Max thinks the chocolate is. This age trend proved remarkably stable after several hundred studies aimed at demonstrating earlier competence (metaanalysis by Wellman et al. 2001).[8]

Our first question is how to represent belief in a mental file system. We can think of two basic approaches. One way is to represent within each file what people think about the object that the file represents, e.g.,:

(1)   Chocolate file: Max thinks this (the chocolate) is in the drawer.

This approach lacks two attractive and intuitive features that simulation theorists (Heal 1986; Gordon 1986) brought to the philosophy of mind: as folk wisdom has it, we understand others by taking their perspective, and we need to bring our own reasoning abilities to bear on that perspective. To satisfy these requirements we need to represent others' perspective in the same format as our own. Provided that we represent our world in terms of mental files, we need to represent other people's perspective with corresponding files. For this reason it is preferable to represent beliefs with vicarious files indexed to different agents (Recanati 2012 chapters 14 and 15), which had also been used by Perner and Brandl (2005); Perner et al. (2007) for representing belief in the false belief task. Mistaken Max' belief about the whereabouts of his chocolate would then look like this:

---

[7] When in a discourse an object is being individuated or referred to under a particular label the corresponding discourse file is activated. A basic mechanism of entrainment or alignment (Pickering and Garrod 2004) obliges the discourse participants to keep conceptualizing the object under this file in order to keep common ground and to not mix different perspectives. These alignment processes have also been described as "conceptual pacts" (Brennan and Clark 1996)

[8] Starting with Clements and Perner (1994) earlier competence could be demonstrated when children are not asked a question about the protagonist, but children's looking behavior is registered, i.e., where they look in expectation of where the protagonist will go, or that they look longer when something unexpected happens than when the expected happens (Onishi and Baillargeon 2005). The interpretation of these data is hotly debated, most recently: Baillargeon et al. 2010; Butterfill and Apperly 2013; Helming et al. 2014; Perner and Roessler 2012; Ruffman 2014.

(2)   Max's Chocolate file: This (the chocolate) is in the drawer

By opting for the use of vicarious files we face the problem of how to relate the vicarious files to the regular files. Here, again, two options come to mind. One option was taken by Perner et al. (2007), p 486) to characterize early pretence. A vicarious file relates to the regular file in that both are anchored to the same external object. This may be enough to engage in pretend play—an interesting question which we will not pursue here[9]—but it will not suffice for an adequate representation of belief, where one has to understand that the chocolate that Max believes is in the drawer is the very same chocolate that is in the cupboard. To represent sameness of referent the vicarious file representing Max's view must be linked to the regular file representing the chocolate. There is an analogy to alternative naming where the discourse files "animal" and "rabbit" must be linked so that children gain access from each file to the information contained in the other files. Linking in this straight sense is, however, not possible in the case of false belief. Linking implies that the information contained in each file can flow freely to any other file to which it is linked. Such linking would lead to confusion and incoherence in the case of false beliefs since the false information in the vicarious file would be absorbed in the regular file and one would lose track of what is actually the case. Information in vicarious files needs to stay quarantined (Leslie 1987). For this reason Recanati (2012) p 193 suggests the term "horizontal linking" for linking that allows for free information flow between regular files, and "vertical linking" for linking vicarious files to regular files.[10]

Introducing vertical linking raises the question how it relates to horizontal linking. Like horizontal linking it is supposed to establish sameness of reference. Linked files have the same referent. Unlike horizontal linking it does not provide free flow of information, only restricted flow. One way information can travel (Recanati 2012 p 193) is by "ascension" (Gordon 1995) from the vicarious to the regular file by marking the content of the vicarious file (e.g., the chocolate is in the drawer) as believed by the person to whom the vicarious file is indexed (e.g., Max believes the chocolate is in the drawer). Another useful measure would be to have free information flow from the regular file to the vicarious file unless there is evidence that the person to whom the vicarious file is indexed does not have the information. This has the desirable effect that under normal circumstances everything we know about, e.g., Max's chocolate we tend to attribute by default to Max. This does not apply, of course, to information we explicitly know he does not have, i.e., that the chocolate was moved to a new location.

In sum we can state that for a proper understanding of belief one needs the ability to link files that share the same referent. This ability is also at the core of being able to manage alternative naming. This commonality gives us a theory of why the two abilities tend to emerge together in development.

---

[9] Pretending that the Lego block is a bar of chocolate could be represented with a discourse file that has no referent (representing a fictional entity) but being anchored by stipulation (no serious recognitional processes) to the Lego piece as a prop. The discourse file functions as part of a discourse: if I pour water over the chocolate it will be wet (Harris and Kavanaugh 1993), etc. The anchoring serves to direct the actions represented in the files to the props.

[10] Perner et al. (2007) e.g., p 488) dealt with vertical linking by having identifiers for each regular file and vicarious files bearing the same identifier, indicating that they share the same referent. But this method falls short of specifying anything about information flow between files.

8 The Emerging Developmental Theory

Our mental files analysis of alternative naming and false belief understanding aimed at finding a common ability required for success so that we get an explanation of the developmental correlation between these tasks around the age of 4 years. Our analysis suggests that the ability to link mental files is the critical element. Without it children can think of an object differently depending on the label with which it has been individuated but do not realize that if individuated one way (e.g., as a rabbit) it is still also the thing it is if individuated the other way (as an animal). Without it children cannot form the vicarious mental files necessary to capture another person's false belief. The developmental theory emerging from these considerations is that around 4 years children become able to link mental files.

To test this theory we need to find new tasks that require linking of mental files and then test whether children start mastering these tasks as they start to understand false belief and pass the alternative naming task. Based on an earlier version of our theory (Perner et al. 2007) it was predicted that children should have tantamount problems processing identity statements. This prediction was confirmed (Perner et al. 2011). For instance, children were asked to figure out which key of a larger set opened which box. The first key opened the yellow box and was thus marked with a yellow sticker on one side. A seemingly new key was taken from the set, but unknown to the child it was the same key as before but shown to the child with the unmarked side. It opened the green box and consequently was marked with a green sticker. Then the experimenter "noticed" that this green key already had the yellow sticker on its other side. She pointed this out to the child and showed the child that the key had the yellow sticker on one side and the green sticker on its other side and said: "Look at the key! This yellow key is the same as the green key!" Then with the key showing its green side the child was asked whether it opened the yellow box and whether it opened the green box. Children readily agreed that it opened the green box but children who failed the false belief task tended to deny that the key would open the yellow box. This is to be expected by analogy to their problems with alternative naming. Children built two different files, one for the green key containing the information that it opens the green box and another file for the yellow key with the information that it opens the yellow box. When faced with the green side of the key, children think of the key as the green key and have the information available that it opens the green box. Children who cannot link files—according to theory, those that fail false belief—do not have the information available that it also opens the yellow box since this information is stored in the inaccessible file for the yellow key.

9 Outlook

To further assess the breadth of our theory we can look for other tasks that are known to be mastered at the same age and correlate with false belief or alternative naming tasks. In fact a number of tasks labelled "perspective tasks" have been found to correlate with false belief understanding: false signs, visual perspective, homonyms, competition/ sabotage (reviewed by Perner and Roessler 2012). The challenge for our theory is to show that they, indeed, require linking of files for success. That this challenge could be

met is indicated by the fact that all these tasks have been taken to require some awareness of perspective differences. Since perspective differences relate to linking files there is some chance that our analysis might contribute to explaining the developmental correlations between these tasks and understanding belief.

Another promising field of application is the so called "intensionality" tasks (Russell 1987), where children have to understand that a referent object known to a person under a single label is not accessible to that person under another label. For instance, suppose Heinz knows where the rubber die is but does not know that the rubber die is also an eraser. Children below the age of about 6 years correctly acknowledge this lack of knowledge. But when asked where Heinz will look for an eraser, they still think that he will look for one where the die is (Apperly and Robinson 2002). Sprung et al. (2007) showed that the needed understanding emerges at the same time as an understanding of second order beliefs, i.e., what a person believes about his own knowledge, on the grounds that both tasks require understanding of embedded perspectives. This analysis can be firmed up with an analysis in terms of mental file management. Several questions are raised by such an analysis. Of interest are the children between 4 and 6 years who pass the false belief test but still have problems with Apperly and Robinson's test. Since they pass the false belief test we can assume that they do use vicarious files and know what believed facts to record on them. We also know that children create vicarious files if the agent has seen the object or has been told about it, and the same holds true for registering facts on those files (Perner and Wimmer 1988). So in Apperly and Robinson's Heinz scenario children who pass the false belief test should be able to create a vicarious file for Heinz, which represents the die and does not contain the information that the die is also an eraser. This allows them to answer correctly that Heinz does not know that the die is an eraser.

Encoding that the die is not an eraser by omitting to encode on the vicarious die-file that it is an eraser is, however, not the same as not creating a vicarious eraser-file for Heinz. How do children—or we—assign vicarious files to an agent? We have to check whether the agent has or has had informational contact with the object and then assign a copy of our regular files to the agent. In particular, if an agent enjoys perceptual acquaintance with the object we create vicarious files for all our files that are anchored to that object—at least at first. For, if we then register that the agent doesn't know that the die or the eraser is an eraser, we need to delete the eraser file.[11] In Heinz's case, children need to first create two vicarious files, a die-file *and* an eraser-file. Then they can register that Heinz does not know that the object is an eraser by omitting the fact of being an eraser in Heinz's vicarious files. They then have to realize that Heinz, not knowing this fact, cannot have an eraser file and, consequently, delete the initially created eraser file. This step, we surmise, incipient false-belief passers don't yet carry out for a couple of years to come (from 4 to 6). With this conjecture we can explain children's strange tendency to let Heinz look for an eraser where the die is when, at the same time, denying him knowledge about it being an eraser. When they see Heinz looking at the eraser-die children assign a vicarious file for both their regular files. Since Heinz does not know that the object is an eraser, they omit from either of the files

---

[11] This sequence is plausible unless we subscribe to simulation theory. Simulation makes it possible to put oneself into the agent's perspective of ignorance about the die being an eraser and, hence, one would not generate a regular eraser file during simulation and, consequently, not tempted to create a vicarious eraser file.

any mention of the fact that the object is an eraser. Importantly, Heinz evidently can see where the object is, so they leave the copied information about its location in each vicarious file. This allows them to answer the question whether Heinz knows that the die is an eraser correctly with "no". However, by hypothesis, the young children stop here and fail to erase the vicarious eraser-file. Consequently, when asked where Heinz might look for an eraser the object remains a viable option for Heinz according to the children (QED).

An open question remains of how children this age deal with linking vicarious files. Two possibilities suggest themselves: they uncritically copy the linking of their regular files or they simply ignore any links. Recent work by Rakoczy et al. (2015) motivates us to adopt the latter. Children and a puppet as an agent observed an object, e.g., a pen being placed in box 1. Then the puppet left (puppet-ignorant condition). In the puppet's absence the pen was taken from box 1 and children were told that the pen was also a rattle and were shown that it rattled when shook. The object was called a "rattle" and returned to box 1. Then the puppet came back and witnessed how an object, hidden in the experimenter's hands, was removed from box 1, described as "rattle," and placed inside box 2. Then children were asked where the puppet would look for the *pen*. The correct answer is "box 1" since the puppet does not know that the *rattle* is also the *pen*. Under the assumption that children uncritically copy their link between regular files to the puppet's vicarious files this question poses a similar problem as Apperly and Robinson's test question about where Heinz will look for an eraser. For, a correct answer requires children to suppress their own knowledge of the identity of die and eraser and in Rakoczy's task to suppress the link between pen and rattle (Rakoczy et al. must have tacitly made this assumption as they dubbed this task "intensionality"). If, however, children ignore links among vicarious files then this test question simply tests understanding false belief: the puppet doesn't know that the pen is being moved, since it can't see it and is only told about a rattle.

In detail the mental file analysis proceeds as follows: the puppet encounters the pen in the beginning when the child herself only knows it as a pen. This leads the child to copy her regular file, which shows the pen in box 1, to a vicarious file. Then the child learns that the pen is also a rattle, which leads to a regular rattle file anchored to the same object. Finally the puppet witnesses that an object called "rattle", which makes a rattling noise, is invisibly moved from box 1 to box 2. Puppet's linguistic contact with *the rattle* lets children copy their regular rattle file to a vicarious rattle file which registers the move to box 2. Puppet's perceptual contact of hearing the object in the experimenter's cupped hands would also lead to a vicarious copy of the pen file—in analogy to the creation of an eraser file for the rubber die in Apperly and Robinson's Heinz scenario. However, unlike the Heinz scenario, the already existing vicarious pen file showing the pen in box 1 will inhibit the creation of a duplicate. Since the child understands that the puppet doesn't have information about the object being a pen the outdated information about its location will not be changed and result in a false belief about location. Hence when asked where puppet will look for the pen they should correctly answer with "box 1". Indeed, the data show that children who answer a false belief question correctly also answer this question correctly, which is in line with the assumption that children ignore links that exist between their regular files when dealing with vicarious files.

Interestingly, if children ignore links among vicarious files then a control condition should become difficult, in which the puppet witnesses everything the child does (puppet knowledgeable condition, labelled "true belief" by Rakoczy et al.). In this case children who fail to take care of links between vicarious files should give the wrong answer because they will not represent that the puppet knows the identity of pen and rattle by linking puppet's vicarious files. Without linking between vicarious files the verbal information about the transfer of the rattle to box 2 will not be registered on the vicarious pen-file. Hence they should answer wrongly that the puppet will look for the pen in box 1. Indeed, children who passed the false belief test gave about 60 % wrong answers (Rakoczy et al. 2015, Table 1, bottom line, rightmost panel headed "true belief").

With this excursion into intensionality experiments on children we wanted to show that mental file analysis provides a rich tool for cognitive analysis and provides a unified explanation for many hitherto unconnected developmental phenomena. Future promising areas for inclusion are children's understanding of others' interpretation of droodles (discrepant findings by Chandler and Helm 1984 versus Taylor 1988 and Perner and Davies 1991), children's problems with knowledge through inference (Sodian and Wimmer 1987), their well-known problems with ambiguous reference (Robinson and Apperly 2001), naming and labelling (Gordon and Olson 1998), with the cardinal number principle (the number for the last object counted is a label for the last object as well as for the size of the set)—to name a few.

Our theory might also become relevant for research on brain function. In particular, the metacognitive awareness needed for managing the linking of files in cases of co-reference may require computational resources localizable in the brain. This would naturally be important for cognitive neuroscience but would also reflect on the importance of mental file theory. There is some evidence that the inferior parietal lobe in the left hemisphere (left IPL) may play an important role. All tasks that require awareness of (the possibility of) different perspectives show activation in this region. False belief (as compared to photo; Saxe and Kanwisher 2003) vignettes activate this region (also on the right). Other theory of mind tasks that also involve perspective activate the same region in left IPL, whereas theory of mind tasks that do not involve perspective—like emotion recognition from eyes (Baron-Cohen et al. 1999) or intentional action infer-ences (Brunet et al. 2000) activate lower regions of IPL and adjacent regions in the temporal lobe (metaanalysis by Schurz et al. 2014). In contrast to false belief vignettes, false sign vignettes activate IPL only on the left (Perner et al. 2006; Aichhorn et al. 2009). Similarly, visual perspective taking tasks (metaanalysis by Schurz et al. 2013), episodic memory (remember minus know judgments; metaanalysis by Arora, A., Weiss, B., Schurz, M., Aichhorn, M., Wieshofer, & Perner, J., Left Inferior-Parietal Lobe Activity in Perspective Tasks: Identity statements. Frontiers in Human Neuroscience (unpublished)), and identity statements activate IPL only on the left. The available evidence speaks thus for regional specificity (Kanwisher 2010), that all these tasks activate specifically left IPL. Whether evidence can also be gained for functional specificity of this region, i.e., that it is exclusively responsible for perspective processes, remains open. The IPL region has been involved in many processes relating to word semantics but also to mathematics tasks like knowledge of equations (overview e.g., Cabeza et al. 2012), to name but a few. To establish functional specificity we would have to show that all these tasks contain (potential) perspective differences,

which is unlikely as it stands. Nevertheless, the possibility of regional specificity suggests that left IPL is responsible for neural activity that is needed in perspective tasks, all of which depend on awareness of linking coreferential files. The existence of such a brain region speaks for the importance of mental files contributing to our understanding of brain function.

References

Aichhorn, M., J. Perner, B. Weiss, M. Kronbichler, W. Staffen, and G. Ladurner. 2009. Temporo-parietal junction activity in theory-of-mind tasks: falseness, beliefs, or attention. *Journal of Cognitive Neuroscience* 21(6): 1179–1192. doi:10.1162/jocn.2009.21082.

Apperly, I.A., and E. Robinson. 2002. Five-year-olds' handling of reference and description in the domains of language and mental representation. *Journal of Experimental Child Psychology* 83(1): 53–75. doi:10.1016/S0022-0965(02)00102-9.

Baillargeon, R., R.M. Scott, and Z. He. 2010. False-belief understanding in infants. *Trends in Cognitive Sciences* 14(3): 110–118. doi:10.1016/j.tics.2009.12.006.

Baron-Cohen, S., H.A. Ring, S. Wheelwright, E.T. Bullmore, M.J. Brammer, A. Simmons, and S.C. Williams. 1999. Social intelligence in the normal and autistic brain: an fMRI study. *European Journal of Neuroscience* 11: 1891–1898.

Brennan, S.E., and H.H. Clark. 1996. Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22: 1482–1493.

Brunet, E., Y. Sarfati, M.C. Hardy-Baylé, and J. Decety. 2000. A PET investigation of the attribution of intentions with a nonverbal task. *NeuroImage* 11: 157–166.

Butterfill, S.A., and I.A. Apperly. 2013. How to construct a minimal theory of mind. *Mind & Language* 28: 606–637.

Cabeza, R., E. Ciaramelli, and M. Moscovitch. 2012. Cognitive contributions of the ventral parietal cortex: an integrative theoretical account. *Trends in Cognitive Sciences* 16: 338–352. doi:10.1016/j.tics.2012.04.008.

Chandler, M.J., and D. Helm. 1984. Developmental changes in the contribution of shared experience to social role-taking competence. *International Journal of Behavioral Development* 7: 145–156.

Clark, E.V. 1997. Conceptual perspective and lexical choice in acquisition. *Cognition* 64: 1–37.

Clements, W.A., and J. Perner. 1994. Implicit understanding of belief. *Cognitive Development* 9: 377–397.

Doherty, M. 1994. *Metalinguistic understanding and theory of mind*. Unpublished doctoral dissertation, Laboratory of Experimental Psychology, University of Sussex.

Doherty, M.J., and J. Perner. 1998. Metalinguistic awareness and theory of mind: Just two words for the same thing? *Cognitive Development* 13: 279–305.

Flavell, J.H. 1988. The development of children's knowledge about the mind: From cognitive connections to mental representations. In *Developing theories of mind*, ed. J.W. Astington, P.L. Harris, and D.R. Olson, 244–267. Cambridge: Cambridge University Press.

Gordon, R.M. 1986. Folk psychology as simulation. *Mind & Language* 1: 158–171.

Gordon, R.M. 1995. Simulation without introspection or inference from me to you. In *Mental simulation: Evaluations and applications*, ed. M. Davies and T. Stone, 53–67. Oxford: Blackwell.

Gordon, A.C., and D.R. Olson. 1998. The relation between acquisition of a theory of mind and the capacity to hold in mind. *Journal of Experimental Child Psychology* 68: 70–83.

Harris, P.L., and R.D. Kavanaugh. 1993. Young children's understanding of pretence. *Monographs of the Society for Research in Child Development* 58(1): 1–107.

Heal, J. 1986. Replication and functionalism. In *Language, mind, and logic*, ed. J. Butterfield, 135–150. Cambridge: Cambridge University Press.

Heim, I. 2002. File change semantics and the familiarity theory of definiteness. In *Formal semantics: The essential readings*, ed. P. Portner and B.H. Partee, 223–248. Oxford: Blackwell Publishing.

Helming, K.A., B. Strickland, and P. Jacob. 2014. Making sense of early false-belief understanding. *Trends in Cognitive Sciences* 18(4): 167–170.

Kamp, H. 1990. Prolegomena to a structural account of belief and other attitudes. In *Propositional attitudes: The role of content in logic, language and mind*, ed. C.A. Anderson, 27–90. Stanford: Center for study of language and information, Lecture Notes Series.

Kamp, H., and U. Reyle. 1993. *From discourse to logic. introduction to model theoretic semantics of natural language, formal logic and discourse representation theory*. Dordrecht: Kluwer Academic Publishers.

Kanwisher, N. 2010. Functional specificity in the human brain: a window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences of the United States of America* 107(25): 11163–11170. doi:10.1073/pnas.1005062107.

Karttunen, L. 1976. Discourse referents. In *Notes from the linguistic underground* (Syntax and semantics, vol. 7), ed. J. McCawley, 363–385. Waltham, MA: Academic Press.

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

Markman, E.M. 1989. *Categorization and naming in children: Problems of induction*. Cambridge: A Bradford Book.

Millikan, R.G. 2000. *On clear and confused ideas. an essay about substance concepts*. Cambridge: Cambridge University Press.

Onishi, K.H., and R. Baillargeon. 2005. Do 15-month-old infants understand false beliefs? *Science* 308: 255–258.

Perner, J., and J. Brandl. 2005. File change semantics for preschoolers: Alternative naming and belief understanding. *Interaction Studies* 6: 483–501.

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

Perner, J., and J. Roessler. 2012. From infants' to children's appreciation of belief. *Trends in Cognitive Sciences* 16(10): 519–525. doi:10.1016/j.tics.2012.08.004.

Perner, J., and H. Wimmer. 1988. Misinformation and unexpected change: testing the development of epistemic-state attribution. *Psychological Research* 50: 191–197. doi:10.1007/BF00310181.

Perner, J., S. Stummer, M. Sprung, and M.J. Doherty. 2002. Theory of mind finds its Piagetian Perspective: Why alternative naming comes with understanding belief. *Cognitive Development* 17: 1451–1472.

Perner, J., M. Aichhorn, M. Kronbichler, W. Staffen, and G. Ladurner. 2006. Thinking of mental and other representations: The roles of left and right temporo-parietal junction. *Social Neuroscience* 1: 245–259. doi:10.1080/17470910600989896.

Perner, J., B. Rendl, and A. Garnham. 2007. Objects of desire, thought, and reality: Problems of anchoring discourse referents in development. *Mind & Language* 22: 475–513.

Perner, J., M.C. Mauer, and M. Hildenbrand. 2011. Identity: Key to children's understanding of belief. *Science (New York, N.Y.)* 333(6041): 474–477. doi:10.1126/science.1201216.

Perry, J. 2002. *Identity, personal identity and the self*. Cambridge: Hackett Publishing.

Pickering, M.J., and S. Garrod. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences* 27(02): 169–190.

Pylyshyn, Z.W. 2007. *Things and places: How the mind connects with the world*. Cambridge: The MIT Press.

Rakoczy, H., D. Bergfeld, I. Schwarz, and E. Fizke. 2015. Explicit theory of mind is even more unified than previously assumed: Belief ascription and understanding aspectuality emerge together in development. *Child Development*. doi:10.1111/cdev.12311.

Recanati, F. 2012. *Mental files*. Oxford: Oxford University Press.

Robinson, E.J., and I.A. Apperly. 2001. Children's difficulties with partial representations in ambiguous messages and referentially opaque contexts. *Cognitive Development* 16: 595–615.

Ruffman, T. 2014. To belief or not belief: Children's theory of mind. *Developmental Review* 34(3): 265–293. doi:10.1016/j.dr.2014.04.001.

Russell, J. 1987. 'Can we say …?' children's understanding of intensionality. *Cognition* 25: 289–308.

Saxe, R., and N. Kanwisher. 2003. People thinking about thinking people: The role of the temporo-parietal junction in "theory of mind.". *NeuroImage* 19: 1835–1842. doi:10.1016/S1053-8119(03)00230-1.

Schurz, M., M. Aichhorn, A. Martin, and J. Perner. 2013. Common brain areas engaged in false belief reasoning and visual perspective taking: a meta-analysis of functional brain imaging studies. *Frontiers in Human Neuroscience* 7: 712. doi:10.3389/fnhum.2013.00712.

Schurz, M., J. Radua, M. Aichhorn, F. Richlan, and J. Perner. 2014. Fractionating theory of mind: A meta-analysis of functional brain imaging studies. *Neuroscience and Biobehavioral Reviews* 42C: 9–34. doi:10.1016/j.neubiorev.2014.01.009.

Sodian, B., and H. Wimmer. 1987. Children's understanding of inference as a source of knowledge. *Child Development* 58: 424–433.

Sprung, M., J. Perner, and P. Mitchell. 2007. Opacity and embedded perspectives: Object identity and object properties. *Mind & Language* 22: 215–245.

Stummer, S. 2001. Sag es anders - ToM! von einem metalinguisitschem bewusstsein zu alternativen erklärungsansätzen. Dissertation, University of Salzburg.

Taylor, M. 1988. Conceptual perspective taking: Children's ability to distinguish what they know from what they see. *Child Development* 59: 703–718.

Tomasello, M. 1999. *The cultural origins of human cognition*. Cambridge: Harvard University Press.

Wellman, H.M., D. Cross, and J. Watson. 2001. Meta-analysis of theory of mind development: the truth about false belief. *Child Development* 72: 655–684. doi:10.1111/1467-8624.00304.

Wimmer, H., and J. Perner. 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition* 13: 103–128.