Supplement to ‘Factive Theory of Mind’

Jonathan Phillips & Aaron Norby

(Main article published in *Mind & Language*)

In the article that this supplement accompanies, we laid out why the long-standing emphasis on false beliefs has led us down the wrong track: it makes the capacity for non-factive representations essential when testing for theory of mind. This approach confounds the kind of content that can be represented with the capacity to represent another agent’s understanding of the world. In place of false beliefs, we argued that a more principled set of criteria consists of (1) tracking (the basic criterion used to test for representation throughout cognitive science), and (2) the ability to keep the tracked representation functionally separate from one’s own understanding of the world. Employing these criteria, we developed the diverse-knowledge task. This task provides a way of testing for the comparatively simple capacity for *factive theory of mind*, while additionally allowing for simple modifications that make it possible to instead test for the ability to represent altercentric ignorance or non-factive content.

We supplement the main article in four primary ways. In §I, we expand on why we think the approach we’ve argued for is a better way of testing for theory of mind from both a theoretical and methodological perspective. We then apply our framework and criteria to a number of existing cases of behavior that may (or may not) provide evidence for theory of mind. In §II, we consider how the framework we’ve outlined provides a principled way for

deciding when a pattern of behavior does not provide evidence for a capacity for theory of mind. Next, in §III, we consider two examples of ways that existing theory of mind research falls short. Specifically, in §III.a we reconsider infant studies involving misleading appearances (e.g., Song and Baillargeon, 2008), and in §III.b we take up interference effects (e.g., Kovács et al., 2010; Samson et al., 2010). In both cases, we argue that these paradigms do not provide clear evidence for theory of mind despite seeming to involve false beliefs. Finally, in §IV, we consider two examples of places where we may have missed good evidence for theory of mind because it didn't involve false beliefs. In §IV.a we consider the evidence from helping paradigms with infants (e.g., Liszkowski et al., 2008), and in §IV.b we consider evidence from monkey theory of mind experiments (e.g., Santos et al., 2006). In both cases, we detail how our framework provides straightforward ways of augmenting these paradigms to ensure that they provide clear evidence for *factive* theory of mind.

I How to test for theory of mind

One of the most important features of the approach we've been arguing for is that it moves away (theoretically and methodologically) from paradigms where a single response is supposed to provide sufficient evidence for a complete capacity for theory of mind.

Part of the point here is theoretical. Unlike the background picture that has often been accepted in theory of mind research, we don't think an ability for theory of mind should be understood as an all-or-nothing capacity to take a certain kind of rationalizing stance when explaining behavior (e.g., Dennett, 1989; Davidson, 1990). Rather, we've argued for a framework on which the standards for representing other minds should be the same as

the standards for other kinds of mental representation. On this picture, the extent to which one has a theory of mind capacity will critically depend on whether or not one can *track* others’ understanding of the world (in exactly the same way that the extent to which one has the capacity to represent number depends on whether one can actually track number specifically). Similarly, the *kind* of theory of mind one will have simply depends on what kinds of mental states one can track (much like the *kind* of numerical capacity one has depends on what kinds of number one can track, e.g., ANS vs. symbolic number representation). A natural consequence of this view is that a single response within a single test or paradigm is just not going to be a good test of what one can track and thus will not be a good way to get a clear picture of the subject’s capacity for theory of mind.

The other part of the point we’re making is empirical. The trouble with relying on any one particular test is that for any given response, there will *always* be some behavioral rule that could in principle explain success on that task Penn and Povinelli (2007).¹ Each of these potential behavioral rules is, in effect, pointing to a way in which a particular pattern of behavior does (or is likely to) covary with the relevant aspect of the agent’s understanding of the situation. Thus in principle, it’s possible that the observed response in the task could have either resulted from participants being sensitive to or tracking that particular pattern of behavior rather than actually tracking the agent’s understanding of the situation.²

¹In fact, given that (1) there will be an indefinitely large number of ways of understanding the particular behavior exhibited by the agent and (2) behavioral rules are learned over the way in which behaviors are understood (not the way that behaviors actually are), there will always an indefinitely large set of behavioral rules that could in principle explain the success.

²Moreover, one need not only be concerned about simple behavioral rules, but also more sophisticated (though still completely nonmentalistic) representations that could explain the behavior (see, e.g., Butterfill and Apperly, 2013; Martin and Santos, 2016; Burge, 2018; Perner and Roessler, 2012).

Maybe we're a little idealistic, but we don't think these facts provide any real reason for despair; it just provides reason to be careful. Each behavioral rule or otherwise nonmentalistic explanation is itself a theory that makes predictions across a wide range of cases. Like any other theory, these alternatives can be tested both directly (by asking whether the same response is observed in cases that don't even potentially involve the representation of mental states) and indirectly (by modifying the original theory of mind paradigm so that the nonmentalistic rule no longer applies but the mental state representation remains and then asking whether the same response is observed). To the extent that direct tests do not provide independent evidence for the proposed alternative theory or the indirect tests undermine the original evidence for the proposed alternative explanation, then these alternative theories can be discounted.³

Taking these points together, the general lesson should be clear: if we are ever going to find sufficient empirical evidence for what a given subject's capacity for theory of mind is, it's not going to be because a subject exhibited a particular response on a particular test; it's going to be because attributing a theory of mind ability is a more parsimonious way of explaining the subject's pattern of responses across different cases. Such an approach will sound immediately familiar to anyone who has engaged in theory or model comparison more generally, and while the underlying proposal here isn't specific in any way to studying theory of mind, the particular way of testing between the relevant alternative explanations is. So what we now want to do in the rest of this paper is turn to a number of instructive examples of

³There are a few important examples where this kind of work is already being done. For example, some non-mentalistic explanations have been argued to not generalize as widely as they would need to (see, e.g., Carruthers (2013); Perner (2014)), while others (e.g., Santiesteban et al. (2014)) have recently been more directly tested and have been argued to not be supported (e.g., Marshall et al. (2018); Gardner et al. (2018)).

how this work can and should be done.

II When and how to decide that some behavior is *not* evidence for ToM

A good starting place is to acknowledge from the outset that there are many instances of behavior that may appear to reveal a capacity for theory of mind but turn out not to in the end. Having a principled way of showing when and why this is the case is an important consequence of the view we're arguing for. Our strategy isn't all that complex; it will just be to ask whether the thing that is being tracked is actually another's understanding, and then to verify that this understanding is kept separate from one's own.

Here's one nice example where much of the necessary work has already been done. A species of a genus of spider, *Portia*, uses 'aggressive mimicry' in order to deceive and prey upon other spiders. The way they do this is by going on to the edge of the prey spider's web and then trying out a series of vibration patterns. They work through different patterns until they reach one that elicits a response from the prey spider, whereupon *Portia* continues to repeat the successful vibration. Generally, the vibration mimics either the vibrations of a trapped insect or of a potential mate. As a result, *Portia* 'calls' the prey-spider closer and closer, until it is less than a body-length away, at which point *Portia* will attack and eat their prey (Jackson and Wilcox, 1993). They also use wind to 'smokescreen' their approaches on to prey spiders' webs, and plan circuitous routes that allow them to drop into a web from above in order to attack spiders that they must approach from behind in order to subdue (Jackson and Wilcox, 1998). Members of *Portia* also sometimes use a 'monotonous web signal that keeps the victim calm'

while it gets close enough to strike.

All of this is behavior that looks like intentional trickery, and it's not difficult to take the intentional stance and start to explain what the spider is doing by saying that it's trying to make the prey-spider think that there's a prey-insect or mate in the web and that it starts out by looking for a vibration that will induce this false belief in the prey spider. Indeed, one could even initially be tempted even take this as preliminary evidence that *Portia* can pass a false-belief task: it expects the prey spider to search where it falsely believes a trapped insect to be, and not elsewhere.

Of course, most researchers do not take this behavior to demonstrate theory of mind ability. Still, some may be tempted to think that a view like ours might be unable to avoid counting this as an instance of theory of mind; after all, isn't this all evidence that *Portia* can track the mental states of their prey? The short answer is *no*, but it's worth explaining why.

The problem is just that all of the evidence suggests that *Portia* do not track their prey's understanding of the situation; they instead track aspects of their prey's *behavior*. When *Portia* goes after a victim, it starts out by generating vibrations from a pre-set repertoire and then, when those do not work, it generates a wide array of new vibrations until it finds one that works (Jackson and Wilcox, 1998). More specifically, it zeros in on a vibration pattern by responding to simple feedback from the prey-spider—specifically, movements from the prey spider that appear to be contingent on one of the *Portia*-induced vibrations (Jackson and Wilcox, 1993). Thus, they are actually following a rule such as: 'Cycle through the repertoire of vibration patterns until a prey spider's approach is detected. Once detected, repeat the most recent vibration pattern until the prey spider is within reach.' The rules required to describe the actual range and flexibility of *Portia*'s aggressive

mimicry are much more complicated than this, but the representational resources required, i.e., what it has to track and represent, are not (for a more in-depth review of work on *Portia* that covers the evidence for what *Portia* tracks, see Harland & Jackson, 2000).

In short, the lesson here should once again be a familiar one. Just like any other area of cognition (e.g., work on number representation in infants, Carey 2009; Feigenson et al. 2004), when doing theory of mind research, one has to show that what is being tracked is the property in question and not something that simply covaries with it. In the case of *Portia*, careful work revealed that what was tracked was something that covaried with prey-spiders' understanding of the situation (the prey-spiders' movement), and not the prey spiders' actual understanding of the situation. Thus, this work does not provide evidence that *Portia* have a capacity for theory of mind; rather, it just shows that they have a capacity for representing the contingent movement of their prey. Our account which emphasizes tracking, provides a principled way of understanding how to tell the difference.⁴

III Where and how existing theory of mind research falls short

We've argued that the false belief task is poorly suited to discovering whether or what kind of theory of mind a particular species or subject has. At the

⁴Note that views like that offered by Martin and Santos (2016) or Butterfill and Apperly (2013) or Burge (2018) take a more intermediate position; they suggest that what is tracked may not be simple behavioral patterns but somewhat richer features, like access or encountering relations, registering an object, teleological movement, conative states, and so forth. To the extent that subsequent research shows that some species or subjects are actually tracking these features rather than other agents' understanding of the world; then this would suggest they have the capacity to represent these features, and perhaps not other agents' actual understanding of the situation. Directly testing whether the proposed features are tracked in cases where they do not covary with other agents' mental states is an important area for future research.

same time, the vast majority of research to date has been conducted with this standard in mind. As such, it is worth returning to the existing theory of mind research to point to places where we think the evidence has not yet provided sufficient support for a genuine capacity for theory of mind.

While we go on to argue that the application of the diverse-knowledge task provides reason for thinking that we may have overlooked important instances of genuine theory of mind ability, it's important to note at the outset that this conclusion was not arrived at by simply relaxing the standards for demonstrating an ability for theory of mind. As we go on to show, rather than requiring less evidence, the proposal we are making requires different (and in some ways stronger) evidence to demonstrate a genuine ability for theory of mind. As a result, there are good reasons to doubt that some widely accepted findings actually provide evidence for a genuine capacity for theory of mind.

The key differences in the requirements proposed by the diverse-knowledge task and the false-belief task can be helpfully illustrated by considering a few examples of the research that has previously been taken to demonstrate an ability for theory of mind. We will focus on two different groups of research which, while clearly suggestive, have not yet provided the evidence required to demonstrate a capacity for theory of mind. To be clear, the argument here is not that participants in these studies were *not* employing theory of mind, rather the argument is that these studies are missing a critical piece of evidence required to *demonstrate* that this is the case. Accordingly, we both lay out precisely what these studies lack and also suggest relatively simple modifications of the experimental paradigms used which would allow them to provide the requisite evidence needed to demonstrate theory of mind—that is to show that they were tracking another agent's understanding of

the situation and keeping that representation separate from their own.

III.a Misleading appearances

Demonstrating that a participant is tracking an agent’s understanding of the situation simply requires evidence that the participant is sensitive to information that is represented therein. Demonstrating separation, however, requires more than this; it also requires evidence that this information is kept functionally separate from the participant’s own representation of the world. Such evidence has not always been provided.

Consider, for instance, the intriguing work on infants’ reasoning about others’ perception of misleading information. In one study (Song and Bailargeon, 2008), 14.5 month old infants were first familiarized with a female agent who demonstrated a preference for one toy (a toy doll with blue hair) over another toy (a skunk) by continually reaching for that toy when given a choice. Subsequently, when this agent was no longer in the room, the experimenter introduced two opaque boxes with lids. One of these boxes had a tuft of blue hair that resembled the doll’s attached to the lid of the box (as the experimenter demonstrated). Then, while the agent was still out of the room, the skunk was placed inside the box with the tuft of blue hair, and the doll was placed inside the other box. The critical test was which box the infants expected the agent to reach for when she returned to the room, as measured by looking time. The infants looked longer when the agent reached for the box where the doll actually was (the box without the blue hair on the lid) suggesting that the infants expected the agent to act on the incorrect inference that the doll was in the box with the blue hair attached to the lid.

This evidence, while suggestive, is not sufficient to demonstrate the capacity for theory of mind because it does not demonstrate that the infants

represented the agent’s understanding of the situation as separate from their own. The easiest way to see this is to realize that it is possible that after the boxes were closed, the infants themselves represented the doll as being in the box with the blue hair attached to the lid. If this were the case, then infants would not need to refer to the agent’s understanding of the situation to predict that the agent would reach for the box with the blue hair attached to the lid; this expectation could have been derived by merely referencing their own understanding of the situation. Of course, it would not be difficult to modify the paradigm to demonstrate that these two representations are kept separate. All that would be required would be additional conditions to demonstrate that infants’ own expectations are violated when the doll is taken from the box with the blue hair but not when the skunk is taken from the same box.

Abstracting away from the details of this particular experimental paradigm, it shouldn’t be hard to see how this point generalize to many other tasks, including many instances of the appearance-reality paradigm⁵. The general lesson is that the diverse-knowledge task puts a strong emphasis on the *separateness* of the representations, and evidence for this separateness cannot be had without directly testing the participant’s own representation. When it comes to demonstrating a genuine capacity for theory of mind, the subject’s own perspective is just as important as the other agent’s.

III.b Interference effects

A number of provocative and important lines of research have focused on demonstrating the presence of theory of mind by arguing that the calculation of theory of mind representations can often interfere with participants’

⁵Note that some studies with older children do include a reality check (Hogrefe et al., 1986; Gopnik and Astington, 1988)

responses on non-theory of mind measures. Most often, these paradigms present participants with visual displays that vary whether the participants' own understanding of the depicted situation differs from the understanding of another agent in the display. Then, experimenters typically test participants' responses on measures that are about the scenario itself (rather than measures that are about the other agent's understanding of the scenario). This approach allows researchers to ask whether participants' responses on non-theory-of mind measures reveal the calculation of the other (often irrelevant) agent's understanding of that situation. Evidence along these lines has been taken to provide support for automatic, implicit, and early-emerging theory of mind abilities.

A concrete example may be helpful. Consider the impressive body of work that has demonstrated automatic Level-1 perspective-taking (Samson et al. 2010; Surtees and Apperly 2012; Surtees et al. 2016b,a; see Apperly 2010 for review). In one of the most well-known studies, participants were shown a room with an avatar standing in the middle and facing one of two walls (Samson et al., 2010). On each trial, a number of dots could be displayed both on the wall that the avatar could see and on the wall that was behind the avatar (which the avatar couldn't see). On half of these trials, participants' task was to indicate how many dots the avatar could see, while on the other half of the trials, participants' task was simply to indicate the number of dots that they themselves saw. The striking finding was that when participants were asked to simply report the number of dots they could see, they were slower (and made more errors) if the avatar in the scene saw fewer dots than they did. That is, even though the avatar was strictly irrelevant to the question at hand, participants seemed to be automatically representing how many dots the avatar saw, and this representation was interfering with

their judgment about what they themselves saw. (A similar methodological approach has also been used to provide evidence for automatic *false belief* representation in human adults (Kovács et al., 2010; van der Wel et al., 2014), and for theory of mind representations in human infants as early as 7 months (Kovács et al., 2010).)

Our worry about this kind of research is structural: the evidence provided by this kind of paradigm is not well-suited to demonstrate a capacity for theory of mind because it is not designed to provide evidence for the *separation* between the participants’ representation of the situation and the representation attributed to the other agent. In fact, it’s just the opposite. In the research by Samson et al. (2010), for example, what is most striking is that the results seem to demonstrate a *lack* of separation between the participants’ and the avatar’s representation of the number of dots present. If it weren’t the case that participants were failing to keep these representations separate, then why were participants’ own responses affected by what the avatar saw? Further evidence for a lack of separation between participants’ own and the avatar’s representation comes from the additional finding that participants also make *egocentric* errors. That is, on trials in which participants were asked to indicate how many dots the avatar saw, they were also slower and made more errors when they themselves saw more dots than the avatar did. Because this kind of paradigm structurally relies on demonstrating a lack of separation to provide evidence for theory of mind calculations, we think it undermines its own ability to provide sufficient evidence for a genuine capacity for theory of mind.

A natural response to this line of argument is to think that we’ve really missed the point. After all, isn’t it precisely the lack of separation that *does* provide evidence that participants are representing the other agent’s

understanding of the world? If participants weren't representing the agent's understanding, then how could that understanding be interfering with participants' responses?!⁶ While we understand this reaction, we worry that this way of thinking is going to get us into trouble. We'll try to explain why.

Start by taking a really clear case. Consider an agent who has the ability to compute any other agent's understanding of any situation, but it turns out she simply cannot keep those representations separate from her own understanding. Regardless of how impressive these mental state computations are, if there's no separation, then at the end of the day, the agent will still end up with a single representation of the situation—even if it is one that is partially informed by computing the other agent's understanding. But if that's right, then for all the reasons laid out in §5.1 of the main article, not only will this agent not be able to pass a standard false belief task, she won't even be able to pass simple knowledge-ignorance tasks: her single map will either have to incorrectly include the facts that she knows (but the other agent doesn't), or incorrectly exclude the facts that the agent doesn't know (but she does). Being able to pass even the simplest theory of mind tasks requires that one be able to keep one's own understanding separate from the other agent's understanding. This is not to say that the agent wouldn't be demonstrating an interesting and perhaps impressive ability; it's just to say that it wouldn't be a theory of mind ability.

Return now to the studies we've been considering. If the only evidence these studies provide is evidence for a *lack* of separation, then there is no

⁶To drive the point home, a colleague once helpfully proposed an analogy to the Stroop effect, which demonstrates that participants' representation of the semantic content of the word interferes with their responses to the visual color of the word (and vice versa). The conflict between the semantic content of the color word and the visual color of the word generates more errors or delayed response times for cues like RED but not cues like RED. Extending the analogy to automatic theory of mind, then, we might rhetorically ask: If participants weren't automatically representing the semantic content of the word, then why do we see the interference with responses to the visual color of the word?

way that this evidence could be decisive. One possibility is that the observed pattern of responses is occurring because that the agent does have two different maps but is having trouble attributing them correctly; another is that it is occurring because participants are simply using other agents' to inform their own single map of the world (and they simply are not maintaining a separate map for the other agent). The trouble is that evidence that is solely derived from interference effects could never tell between the two, and so evidence of interference *by itself* just isn't good evidence for a genuine capacity for theory of mind.

Once again, it's not hard to see what would additionally be required to show that the ability being demonstrated is a genuine ability for theory of mind. One needs only to show that participants have in fact separately maintained their own representation of the situation, and that it differs from that of the other agent. One cannot tell this from interference alone.

It's probably worth reiterating: our argument is not that all of the tasks that have shown interference effects absolutely do not involve participants engaging in theory of mind. Rather, it's that this kind of effect does not provide good evidence for it.⁷ To demonstrate a genuine capacity for the-

⁷A separate line of argument has been made that some of these effects do not arise from theory of mind processes at all. One general way to put this empirical worry is that the thing these studies provide *direct* evidence for is a change in participants' responses about their own understanding of the situation, but they don't, however, give us direct evidence that the representation of another agent's understanding of that situation changed participants' responses. Thus a general worry is that some other feature of the experiment is actually responsible for changing participants' responses (see, e.g., Heyes, 2014b). To take one example, in the Samson et al. (2010) paradigm, participants may tend to pay more attention to the dots that the other agent is looking at, and thus may be overly focused on just those dots when responding to prompts about how many dots they saw (a possibility that is especially concerning given that replacing the agent with an arrow that points to where the agent looked can generate a similar pattern of errors, Santiesteban et al. 2014, though see Marshall et al. 2018 and Gardner et al. 2018 for evidence that questions the specific attention-based account offered in Heyes (2014b) and Santiesteban et al. (2014). Or, to take another example, in the Kovács et al. (2010) paradigm, the timing of the events in the visual display co-varied with what the other agent should have believed in way that that led to the relevant differences in adult participants' response times (Phillips et al., 2015).

ory of mind, it is critical not only to show that participants are tracking other agents' understanding of the situation but also that participants' own understanding is kept functionally *separate* from other agents'.

IV Promising places where we may have missed evidence for theory of mind (and how to tell)

Finally, we want to return to a couple of studies that we believe provide potentially promising evidence for *factive* theory of mind. In each case, we consider the best alternative explanation that does not require tracking others' understanding and keeping that separate from your own understanding. We then go on to show how our approach provides simple ways of testing between these competing explanations. These tests are not only helpful in the particular experimental paradigms we consider, but serve to illustrate much more general principles that can be easily extended to other research paradigms.

IV.a How to figure out how helpful infants really are

As described earlier, infants as young as 12 months will help an adult find an object (e.g., a stapler) that moved locations when the adult wasn't watching, but not when the object changed locations in clear view of the adult (Liszkowski et al., 2006, 2008). What's intriguing about this is that, completely setting aside the debate over whether infants represent others' false beliefs (Onishi and Baillargeon, 2005; Southgate and Vernetti, 2014; Heyes, 2014a), these comparatively simple helping studies have the potential to provide evidence of a genuine theory of mind ability without involving the representation of beliefs whatsoever. After all, if infants did not represent

the fact that the adult is ignorant of the location of the object (tracking the agent’s understanding of the situation), and at the same time keep track of where that object really is (thus keeping their own representation of the situation separate), they would likely be unable to succeed on this task. As such, these studies can be thought of as functioning as a modified version of the diverse-knowledge task.

An important alternative interpretation of infants’ success on these tasks is that infants are instead relying on the kind of awareness relations suggested by Martin and Santos (2016). On this alternative proposal, young children only have one map, which includes the actual location of the desired object; in one condition they represent the experimenter as having access to this map (including the place where the object moved) and thus don’t try to update the experimenter’s map by pointing; in the other condition, they don’t attribute a map to the other agent at all, and simply point based on their own understanding of the situation.⁸

While we don’t know of extant evidence that settles this issue, our account provides a principled and easy way of testing if subjects are relying on awareness relations rather than factive theory of mind (see §5.1). In the context of the helping experiments by Liszkowski and colleagues, the key question is whether in the ignorance condition, infants are actually attributing *no* map to the adult, or whether they are selectively representing the adult as *ignorant* of some particular part of the situation involving the stampler. Figuring out which of these is right is not difficult. One simply needs to test how selective the ignorance is. For example, suppose that the en-

⁸We’d like to thank an anonymous reviewer at *M&L* for suggesting this alternative. This interpretation is particularly salient given the recent evidence for a teleological or otherwise nonmentalistic explanation of previous patterns of behavior in helping studies that were thought to provide evidence for false belief understanding at 18 months (see, e.g., Allen 2015 and Priewasser et al. 2017 on the helping task in Buttelmann et al. 2009).

tire experiment began with the infant watching the adult demonstrate a preference for a single object and then placing that object in a closed box. The experiment then proceeds as before, with the adult also displaying an interest in one of two new objects which are then transferred to different locations out of view of the experimenter. If the infant is using factive theory of mind, then when the adult begins searching, they should selectively point to the preferred object that moved (and not to the preferred object that simply stayed in the box). If the infant is not attributing any map, then they should simply point to whichever object they think the adult has the greatest preference for, irrespective of whether or not that object moved without the adult’s knowledge.

The details here are of course specific to the particular experiments by Liszkowski and colleagues, but the point is easily made more general: In any case where success on the task may occur by attributing *no* map rather than selectively attributing ignorance, one simply needs to test whether the relevant agent is still understood to be aware of some *other* part of the situation, which has not changed. If so, then the study provides evidence that the participant is attributing a separate, different map (rather than no map) and thus is demonstrating a genuine capacity for theory of mind. This point is just as relevant to the these infant helping studies as it is of research with non-human primates (e.g., (Martin and Santos, 2014)) or even work on implicit theory of mind representations in adults.

IV.b How to know what non-human primates know you know

Another example of an extant paradigm that captures the essence of the diverse-knowledge task (even if not under that description), is the work on Rhesus macaque theory of mind described earlier (Santos et al., 2006). Sim-

ilar studies have also been done with other primate species (Bräuer et al., 2007; Hare et al., 2006; Melis et al., 2006). Hare et al. (2006) used a conceptually similar experimental paradigm to study chimpanzee’s understanding of what other agents have visual access to. Like the Rhesus macaques, the chimpanzees were able to choose the action that would not alert the experimenter, even on their first attempt, providing evidence for both a calculation of which parts of the situation the experimenter did not understand, and some way of keeping that understanding separate from their own. Moreover, using a slightly different paradigm, researchers also found that chimpanzees chose actions that advantageously manipulated either the auditory and visual evidence the experimenter had access to, and that this behavioral pattern was directly contingent on what the experimenter could see or hear (Melis et al., 2006), providing further evidence for tracking and separation in non-human primates.

Rather than going through these studies in more detail, we want to return to the question of whether there may be some alternative nonmentalist way to explain the pattern of behavior observed. For example, building on the previously discussed proposal by Martin & Santos (2016), could it be that rhesus monkeys in these cases were simply *not* representing and attributing a separate map to the experimenter in one condition, but were in another? Or alternatively, could it be that, as in the case of *Portia*, there are much simpler non-mentalist mechanisms that actually explain the behavior (e.g., a learned negative association with making noise or some gaze-based heuristic)? Using the Santos et al. (2006) study as an example, we’ll consider these possibilities in turn.

Can the pattern be explained by the lack of an attribution of a map (as in Martin and Santos 2016) rather than factive theory of mind? The feature

that makes the Santos et al. (2006) study particularly intriguing is that it is specifically in the condition where the experimenter is not looking (and thus would *not* be attributed a map) that the monkeys *preferentially* go to the silent box rather than the noisy box. When the experimenter looks at the monkey, this preference, if anything, is reversed. This makes it particularly difficult to explain the pattern of behavior on an account where the monkeys are simply not attributing any map to the experimenter in one condition but are in another. After all, if they aren't attributing a map when the agent is looking away, then why do they preferentially go to the silent rather than the noisy box rather than choosing randomly between the two boxes? Even more perplexingly, why would this pattern then go away (or even reverse) when the experimenter *is* looking at them (i.e., when they are supposed to be attributing a map to the experimenter)? Something more than simply not attributing a map is needed.

Perhaps though, it is possible to explain this behavior by proposing a some nonmentalistic behavioral rule. The obvious initial idea would be that the monkeys may have learned that a good heuristic to live by is to not make noise when getting food when a dominant other is present. This is a plausible rule the monkeys may have learned, but it obviously can't be what explains these data, since the monkeys don't exhibit a preference for the silent box when the experimenter is looking directly at them. A revised version that would provide a better fit to the data is that the monkeys may have learned a more complex conditional heuristic—something of the form: don't make noise when taking food if a dominant other is present but *no* direct gaze is detected; else, don't worry about making noise (e.g., in cases where the dominant other is not present, or is present but gaze is detected). In other words, they've learned that there is a connection between making noise and

the occurrence of a dominant others' gaze (or physical interaction) that is only relevant when the dominant other's gaze is not already detected.⁹

While we don't know of any independent evidence in favor of this specific conditional rule, the goal of the framework we've been developing is that it should give us a straightforward way to ask whether the observed behavior is best explained by a behavioral heuristic or by a capacity for factive theory of mind.

According to a factive theory of mind interpretation, the observed behavior critically depends on the fact that the monkey was tracking the experimenter's understanding of the situation—attributing a map that includes both the silent and noisy boxes of food. If no map had been attributed, then the monkeys should not have been sensitive to the fact that opening the noisy box would have alerted the dominant other to the fact that food was being taken. In contrast, the nonmentalistic account suggests that the observed pattern does not depend on the map attributed to the experimenter, but only on the monkey's own understanding of the situation: specifically, they were tracking the fact that one box was noisy and that the experimenter (a dominant other) was present but no gaze was detected. Accordingly, the principled way to test between these two competing explanations is to ask what happens when the monkeys' understanding of the situation differs from the experimenter's (the dominant other's).

One straightforward way to do this is to have a condition in which both the silent and the noisy box are placed in front of an occluder before the experimenter is present, so that only the monkey, but not the experimenter has seen the boxes. In such a case, the heuristic-based explanation predicts

⁹The emphasis on gaze and behavior rather than, e.g., the dominant other *seeing* you take food is important, as representations of *seeing* would themselves be factive mental state representations.

that the same preference for the silent box will be exhibited, as the monkey should again be tracking the fact that one box is noisy and that a dominant other is present but no gaze is detected. In contrast, a factive theory of mind account, predicts that the monkey will no longer be concerned that opening the noisy box will alert the experimenter that food is being taken, since neither the boxes nor the food in them will be represented as part of the experimenter’s map. Accordingly, a factive theory of mind account does not predict a preference for the silent box under these conditions, but the conditional heuristic account does.

While the details here are clearly specific to the setup of the Santos et al. (2006) study, the structural point is again perfectly generalizable. Any account based on a heuristic (or a behavioral rule, a learned association, or any nonmentalistic feature that covaries with the agent’s understanding of the situation) will depend only on features of the subject’s own representation of the situation. Thus these accounts can be tested against factive theory of mind accounts by manipulating whether the other agent’s understanding includes the relevant part of the situation while holding fixed the proposed features that were part of the subject’s own representation. This general point should be useful in determining whether or not many different patterns of behavior in non-human primates are evidence for factive theory of mind (e.g., Hare et al. 2006; Melis et al. 2006).

If, as we suspect will be the case, additional tests along these lines find that these primates are not simply tracking nonmentalistic features, then it’s worth pointing out that in these cases, great apes and even monkeys are succeeding at these tasks by exercising a capacity for theory of mind that is completely *factive*. This possibility is particularly intriguing in light of the fact that there is currently no good evidence to date that monkeys

have any capacity for non-factive representations (Martin and Santos, 2014, 2016), and only mixed evidence in the case of great apes (Kaminski et al., 2008; Krachun et al., 2009; Krupenye et al., 2016). Thus, despite the paucity of clear evidence for non-factive representations, these studies may provide examples of the capacity for both tracking what another agent understands and keeping that representation functionally separate from their own understanding of the world. That is, they may provide examples of the capacity for factive theory of mind in the absence of any capacity for non-factive theory of mind.

References

- Allen, J. W. (2015). How to help: Can more active behavioral measures help transcend the infant false-belief debate? *New Ideas in Psychology*, 39:63–72.
- Apperly, I. (2010). *Mindreaders: the cognitive basis of" theory of mind"*. Psychology Press.
- Bräuer, J., Call, J., and Tomasello, M. (2007). Chimpanzees really know what others can see in a competitive situation. *Animal cognition*, 10(4):439–448.
- Burge, T. (2018). Do infants and nonhuman animals attribute mental states? *Psychological Review*, 125(3):409–434.
- Buttelmann, D., Carpenter, M., and Tomasello, M. (2009). Eighteen-month-old infants show false belief understanding in an active helping paradigm. *Cognition*, 112(2):337–342.

- Butterfill, S. A. and Apperly, I. A. (2013). How to construct a minimal theory of mind. *Mind & Language*, 28(5):606–637.
- Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.
- Carruthers, P. (2013). Mindreading in infancy. *Mind & Language*, 28(2):141–172.
- Davidson, D. (1990). The emergence of thought. *Erkenntnis*, 51:7–17.
- Dennett, D. C. (1989). *The intentional stance*. MIT press.
- Feigenson, L., Dehaene, S., and Spelke, E. (2004). Core systems of number. *Trends in cognitive sciences*, 8(7):307–314.
- Gardner, M. R., Bileviciute, A. P., and Edmonds, C. J. (2018). Implicit mentalising during level-1 visual perspective-taking indicated by dissociation with attention orienting. *Vision*, 2(1):3.
- Gopnik, A. and Astington, J. W. (1988). Children’s understanding of representational change and its relation to the understanding of false belief and the appearance-reality distinction. *Child development*, pages 26–37.
- Hare, B., Call, J., and Tomasello, M. (2006). Chimpanzees deceive a human competitor by hiding. *Cognition*, 101:495–514.
- Harland, D. P. and Jackson, R. R. (2000). Eight-legged cats and how they see: A review of recent research on jumping spiders (araneae: Salticidae). *Cimbebasia*, 16:231–240.
- Heyes, C. (2014a). False belief in infancy: a fresh look. *Developmental science*, 17(5):647–659.

- Heyes, C. (2014b). Submentalizing: I am not really reading your mind. *Perspectives on Psychological Science*, 9(2):131–143.
- Hogrefe, G.-J., Wimmer, H., and Perner, J. (1986). Ignorance versus false belief: A developmental lag in attribution of epistemic states. *Child development*, pages 567–582.
- Jackson, R. R. and Wilcox, R. S. (1993). Spider flexibly chooses aggressive mimicry signals for different prey by trial and error. *Behaviour*, 127(1/2):21–36.
- Jackson, R. R. and Wilcox, R. S. (1998). Spider-eating spiders. *American Scientist*, 86(4):350–7.
- Kaminski, J., Call, J., and Tomasello, M. (2008). Chimpanzees know what others know, but not what they believe. *Cognition*, 109(2):224–234.
- Kovács, A. M., Téglás, E., and Endress, A. D. (2010). The social sense: Susceptibility to others’ beliefs in human infants and adults. *Science*, 330:1830–34.
- Krachun, C., Carpenter, M., Call, J., and Tomasello, M. (2009). A competitive nonverbal false belief task for children and apes. *Developmental Science*, 12(4):521–535.
- Krupenye, C., Kano, F., Hirata, S., Call, J., and Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354(6308):110–114.
- Liszkowski, U., Carpenter, M., Striano, T., and Tomasello, M. (2006). 12- and 18-month-olds point to provide information for others. *Journal of Cognition and Development*, 7(2):173–187.

- Liszkowski, U., Carpenter, M., and Tomasello, M. (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition*, 108(3):732–739.
- Marshall, J., Gollwitzer, A., and Santos, L. R. (2018). Does altercentric interference rely on mentalizing?: Results from two level-1 perspective-taking tasks. *PloS one*, 13(3):e0194101.
- Martin, A. and Santos, L. R. (2014). The origins of belief representation: Monkeys fail to automatically represent others’ beliefs. *Cognition*, 130:300–8.
- Martin, A. and Santos, L. R. (2016). What cognitive representations support primate theory of mind? *Trends in Cognitive Sciences*, 20(5):375–382.
- Melis, A. P., Call, J., and Tomasello, M. (2006). Chimpanzees (pan troglodytes) conceal visual and auditory information from others. *Journal of Comparative Psychology*, 120(2):154.
- Onishi, K. H. and Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, 308:255–8.
- Penn, D. C. and Povinelli, D. J. (2007). On the lack of evidence that non-human animals possess anything remotely resembling a ‘theory of mind’. *Philosophical Transactions of the Royal Society B*, 362:731–44.
- Perner, J. (2014). Commentary on ted ruffman’s “belief or not belief:...”. *Developmental Review*, 34(3):294–299.
- Perner, J. and Roessler, J. (2012). From infants’ to children’s appreciation of belief. *Trends in cognitive sciences*, 16(10):519–525.

- Phillips, J., Ong, D. C., Surtees, A. D., Xin, Y., Williams, S., Saxe, R., and Frank, M. C. (2015). A second look at automatic theory of mind: Reconsidering kovács, téglás, and endress (2010). *Psychological Science*, 26(9):1353–1367.
- Priewasser, B., Rafetseder, E., Gargitter, C., and Perner, J. (2017). Helping as an early indicator of a theory of mind: Mentalism or teleology? *Cognitive Development*.
- Samson, D., Apperly, I. A., Braithwaite, J. J., Andrews, B. J., and Scott, S. E. B. (2010). Seeing it their way: Evidence for rapid and involuntary computations of what other people see. *Journal of Experimental Psychology: Human Perception and Performance*, 36(5):1255–66.
- Santesteban, I., Catmur, C., Hopkins, S. C., Bird, G., and Heyes, C. (2014). Avatars and arrows: Implicit mentalizing or domain-general processing? *Journal of Experimental Psychology: Human Perception and Performance*, 40(3):929.
- Santos, L. R., Nissen, A. G., and Ferrugia, J. A. (2006). Rhesus monkeys, *Macaca mulatta*, know what others can and cannot hear. *Animal Behaviour*, 71(5):1175–1181.
- Song, H. and Baillargeon, R. (2008). Infants’ reasoning about others’ false perceptions. *Developmental Psychology*, 44(6):1789–95.
- Southgate, V. and Verneti, A. (2014). Belief-based action prediction in preverbal infants. *Cognition*, 130:1–10.
- Surtees, A., Apperly, I., and Samson, D. (2016a). I’ve got your number: Spontaneous perspective-taking in an interactive task. *Cognition*, 150:43–52.

- Surtees, A., Samson, D., and Apperly, I. (2016b). Unintentional perspective-taking calculates whether something is seen, but not how it is seen. *Cognition*, 148:97–105.
- Surtees, A. D. and Apperly, I. A. (2012). Egocentrism and automatic perspective taking in children and adults. *Child Development*, 83(2):452–60.
- van der Wel, R. P., Sebanz, N., and Knoblich, G. (2014). Do people automatically track others’ beliefs? evidence from a continuous measure. *Cognition*, 130(1):128–133.