

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Evidence for evaluations of knowledge prior to belief

Permalink

<https://escholarship.org/uc/item/1dp0j198>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Phillips, Jonathan

Knobe, Joshua

Strickland, Brent

et al.

Publication Date

2018

Evidence for evaluations of knowledge prior to belief

Jonathan Phillips¹ (phillips01@g.harvard.edu), Joshua Knobe² (joshua.knobe@yale.edu)
Brent Strickland³ (stricklandbrent@gmail.com), Pauline Armory³ (pauline.armory@gmail.com)
& Fiery Cushman¹ (cushman@fas.harvard.edu)

¹Department of Psychology, Harvard University, 33 Kirkland St. Cambridge, MA, 02138

²Program in Cognitive Science, Yale University, 2 Hillhouse Ave. New Haven, CT, 06520-8205

³Département d'Études Cognitives, Institut Jean Nicod, 29, rue d'Ulm, 75230 Paris cedex 05

Abstract

We investigate the relationship between evaluations of knowledge and belief in human adult theory of mind, and provide evidence that evaluations of knowledge are made without prior evaluations of belief. Our studies find that (1) people can accurately evaluate others' knowledge before they evaluate their beliefs; (2) this pattern cannot be explained by pragmatic differences; (3) it occurs cross-linguistically and unlikely to be accounted for by differences in word frequency, and (4) it also generalizes to the larger class of factive and non-factive attitudes (to which knowledge and belief respectively belong). Together, these studies demonstrate that human adults can ascribe knowledge without first ascribing a belief state. More generally, they lend support to the view that knowledge representations are a distinctive and basic way in which we make sense of others' minds.

Keywords: knowledge; belief; theory of mind; factive attitudes; non-factive attitudes; False Belief; knowledge first

Introduction

Ordinarily, we say that a person 'knows' some proposition only if she both believes it and it is true. For instance, it sounds wrong to say *Jane knows that it is raining* if there isn't a cloud in the sky; likewise if it really is raining, but Jane believes that it is snowing. This relationship is often captured by the idea that knowledge entails belief (Gettier, 1963; Lehrer, 1968; Williamson, 2002).¹ Given this relationship between knowledge and belief, an obvious hypothesis is that in evaluating whether a person knows something, one must first determine whether they believe that thing. How else would the attribution of knowledge be restricted to cases of belief? While this first hypothesis might seem highly intuitive, there is also some evidence that evaluations of knowledge do not depend on prior evaluations of belief. For example, the capacity for knowledge representation appears to have evolved first and to emerge earlier in human development. Our aim is to contribute to this debate by investigating a new question: Can adults judge what somebody knows faster than they can judge what somebody believes?

Two opposing views of belief and knowledge

Previous research on theory of mind provides support for two conflicting ways of understanding the relationship between knowledge and belief. One view suggests that knowledge depends on belief, and thus belief must be more basic than

knowledge; the other suggests that knowledge does not depend on belief, and thus need not be more basic.

Belief before knowledge Within theory of mind research, a standard assumption is that the capacity to represent others' beliefs is among the most basic ways in which we understand others' minds. This assumption appeared early in theoretical work on theory of mind (Dennett, 1987), and continues today in the standard assumption that an ability to represent beliefs is necessary for a "genuine" capacity for theory of mind (Baron-Cohen, Leslie, & Frith, 1985; Call & Tomasello, 2008; Leslie, 1987).

A natural extension of this view is that representations of other mental states, e.g., what others *hope for*, or *suppose*, are often composed over elements including the representation of belief (Snyder et al., 1991). On this view, knowledge requires belief, but it also extends beyond it. For example, a standard philosophical view is that knowing that *p* can be understood as an instance of (1) having some belief that *p*, (2) *p* being true, (3) being justified in having the belief that *p*, and (4) satisfying some number of additional criteria (see Ichikawa & Steup, 2016, for a helpful overview of the kinds of criteria that have been proposed). Consistent with this picture, empirical research has provided evidence that the capacity to represent others' beliefs emerges extremely early in human development (Onishi & Baillargeon, 2005), and that even non-human primates may be able to beliefs (Krupenye, Kano, Hirata, Call, & Tomasello, 2016). Taken together, this research suggests that the ability to represent beliefs may be the among the most basic components of theory of mind.

This picture of the relationship between knowledge and belief predicts that human adults should be able to correctly evaluate others' beliefs before they are able to correctly evaluate others' knowledge. After all, attributions of beliefs should be comparatively simple; moreover, interpreted strictly, this picture posits belief representation as an input to computations of knowledge.

Knowledge before belief In contrast to this picture, there is some evidence that representations of knowledge do not actually depend on prior representations of belief. Studies of language development show that children begin productively using the term 'know' before 'believe' or 'think' (Bretherton & Beeghly, 1982). They can also successfully evaluate others' knowledge states before they succeed in evaluating their belief states (see, e.g., Mar, Tackett, & Moore, 2010). Similar

¹In these experiments are only concerned with propositional knowledge and not practical knowledge or knowledge-*wh*. For challenges to the view that knowledge entails belief, see, e.g., Radford, 1966; Myers-Schulz & Schwitzgebel, 2013.

evidence extends beyond the specific concepts of knowledge and belief. More generally, one of the central differences between knowledge and belief is that representations of knowledge are *factive*, meaning that one cannot represent others as *knowing* something false (while one can represent others as *believing* something false).² Within theory of mind research, there is an enormous amount of work suggesting that factive attitude representations (e.g., seeing, hearing, being aware) may be simpler or more basic than non-factive representations. This is true for example, in developmental work, where there is clear evidence for the representation of others' factive mental states substantially before there is any evidence for non-factive mental state representations (i.e., 6 months vs. 8-12 months) (Luo & Johnson, 2009; Vouloumanos, Martin, & Onishi, 2014). Additionally, within non-human primate research, there is overwhelming evidence that a range of non-human primates can represent factive attitudes (Santos, Nissen, & Ferrugia, 2006), and even some convincing evidence that they fail to represent beliefs (Martin & Santos, 2014, 2016). As a whole, this research paints a picture on which representations of knowledge may not depend on beliefs, since knowledge representations can occur without the capacity to represent belief.

On this picture of the relationship between knowledge and belief, it need not be the case that human adults are able to correctly evaluate others' beliefs before they are able to correctly evaluate others' knowledge. In fact, adults may actually be faster to evaluate knowledge than belief.

The present studies

A series of four experiments test these two opposing predictions about human adults' response times in evaluating others' knowledge and beliefs. Experiments 1a-b provide a straightforward test and find evidence that knowledge evaluations are faster than belief evaluations. Experiment 2 asks whether this difference can be accounted for by differences in the pragmatics of statements about knowledge versus belief, and finds that it cannot. Experiment 3 then tests whether this pattern occurs cross-linguistically and provides evidence that it does. Finally, Experiment 4 demonstrates that this difference in response times generalizes to the larger class of factive and non-factive attitudes, to which knowledge and belief respectively belong. Together, these experiments suggest that in determining what others know, human adults do not first determine what they believe.

Experiment 1a-b

Methods

Participants. In Experiment 1a, 200 participants ($M_{age} = 32.76$, $SD_{age} = 12.67$; 108 females) were recruited through a psychology based website

²Within linguistics, the distinction between factive and non-factive attitudes is roughly that factive attitude ascription, e.g., 'S knows that p ', presuppose that the complement p , while non-factive attitude ascriptions, e.g., 'S believes that p ', do not presuppose p (Kiparsky & Kiparsky, 1970).

(<http://www.moralsensetest.com/>). Experiment 1b was an exact replication involving 501 new participants recruited through Amazon Mechanical Turk.

Stimuli and procedure. Participants began by completing a demographic questionnaire and two practice trials that familiarized them with the task. Participants then completed twenty-four trials in which they read short vignettes about agents and decided whether a sentence about the story was true or false. Participants were instructed to indicate their responses as quickly as possible by pressing one of two keys on their keyboard. Twelve of the trials were distracter trials that were included to prevent participants from anticipating the critical sentences. In the other twelve trials, participants read vignettes that described an agent as either having a true belief about some proposition p (**True Belief**), being ignorant of p (**Ignorance**), or believing some proposition q that was both false and inconsistent with p (**Different Belief**), such as:

True Belief: Mira looks at the night sky with her telescope. She owns the most accurate books on the locations of the different planets throughout the year. Mira reads in her astronomy books that she can see Neptune through her telescope, and she waits until it's dark enough outside. She points her telescope towards the coordinates that her books specify for Neptune, and sees a bright dot in the middle of the sky. That bright dot is Neptune. She is excited that she found the planet she was looking for so easily.

Ignorance: Mira likes looking at the night sky with her telescope. She owns the most accurate books on the locations of the different planets throughout the year. It is night and Mira decides not to read her astronomy books and instead just look through her telescope. Ignoring her book, she sets up her telescope and points it towards a group of dots that catch her attention. She looks into the telescope and she sees a bright dot in the middle of the sky. That bright dot is actually Neptune.

Different Belief: Mira likes looking at the night sky with her telescope. She owns the most accurate books on the locations of the different planets throughout the year. It is night and Mira reads in her astronomy books that she can see Mercury through her telescope. Misreading her book, she sets up her telescope and points it towards the coordinates that her books specify for Neptune, which also happens to be in the sky. She looks into the telescope and she sees a bright dot in the middle of the sky. That bright dot is actually Neptune.

On these trials, participants evaluated the truth of a sentence about knowledge, as in (1), or belief, as in (2). We used 'thinks' instead of 'believes' to equate for frequency/length.

- (1) Mira knows she is looking at Neptune.
- (2) Mira thinks she is looking at Neptune.

Analysis approach

In all of our studies, response times for trials on which participants correctly assessed the truth of the knowledge and belief statements were analyzed with linear mixed effects models using the lme4 package in R (Bates, Maechler, Bolker, Walker, et al., 2014), with both participants and scenario included as random factors. To determine the significance of particular effects, we compare a model that includes the relevant term in question (as well as all other factors not being investigated) with a model that differs only in not including that term. The effect is taken to be significant if the fit of the two models differs significantly (Barr, Levy, Scheepers, & Tily, 2013). Participants were excluded from the analysis if they answered less than 60% of the questions correctly or if their mean response time was less than 1000ms or greater than 4000ms. We additionally excluded all trials on which 1000ms > response time > 4500ms. Fewer than 15% of participants were excluded from any study.

Experiment 1a results

The overall analysis of participants response times revealed no main effect of Belief Condition, $\chi^2(2) = 1.445$, $p = .486$, and no Belief Condition \times Ascription Type interaction, $\chi^2(2) = 1.615$, $p = .446$.³ However, there was a significant main effect of Ascription Type, $\chi^2(1) = 22.382$, $p < .001$, such that participants were faster to correctly assess the truth of statements about whether the agent knows something ($M = 2814.45$, $SD = 935.82$) than statements about whether the agent thinks something ($M = 2991.30$, $SD = 986.34$).

Experiment 1b results

The analysis of participants response times in a well-powered replication of Experiment 1a again revealed no main effect of Belief Condition, $\chi^2(2) = 1.716$, $p = .424$, but did reveal a main effect of Ascription Type, $\chi^2(1) = 27.687$, $p < .001$, such that participants were faster to correctly assess the truth of statements about whether the agent knows something ($M = 2485.92$, $SD = 655.13$) than statements about whether the agent thinks something ($M = 2545.73$, $SD = 660.12$). Additionally, there was a significant Belief Condition \times Ascription Type interaction, $\chi^2(2) = 23.85$, $p < .001$. We next combined the data from both experiments to ask whether the interaction was significant across all of the data.

Combined analysis and discussion

The combined analysis again showed no main effect of Belief Condition, $\chi^2(2) = 2.215$, $p = .330$, a significant main effect of Ascription Type, $\chi^2(1) = 48.418$, $p < .001$, such that participants were overall faster to correctly assess the truth of statements about whether the agent knows something ($M =$

³The fixed and random effects structure for the full model was specified as: $\text{response.time} \sim \text{belief.condition} * \text{ascription.type} + (1|\text{scenario}) + (\text{belief.condition} * \text{ascription.type}|\text{subj})$. We were not able to include random slopes for the scenario because the crossed nature of the random effects in our experiment prevented convergence. We employ a similar approach throughout our studies.




Figure 1: Mean response time across Experiment 1a-b for evaluations of knowledge ascriptions (dark bars) and belief ascriptions (light bars) as a function of Belief Condition. Error bars depict ± 1 SEM.

2508.20, $SD = 653.03$) than statements about whether the agent thinks something ($M = 2587.25$, $SD = 661.14$). Moreover, there was a significant Belief Condition \times Ascription Type interaction, $\chi^2(2) = 12.352$, $p = .002$. Planned pairwise comparisons revealed that participants' response times only differed in the True Belief, $t(450) = 2.525$, $p = .012$, $d = 0.121$, and Ignorance, $t(436) = 5.331$, $p < .001$, $d = 0.270$ conditions, but not in the Different Belief condition, $t(442) = 1.959$, $p > .05$, $d = 0.097$ (see Figure 1).

Our results show clearly that participants can correctly evaluate knowledge ascriptions before belief ascriptions. We will return to why this effect is attenuated in the **Different Belief** condition in the following experiment.

Experiment 2

Before continuing, we want to directly address the concern that the observed difference in response times could have arisen from an unintended pragmatic implicature. More concretely, there is reason to think that there may be something pragmatically odd about belief ascriptions (as in (2)) but not knowledge ascriptions (as in (1)) in cases where the agent seems to actually know the relevant proposition (as in **True Belief**), see e.g., Heim, 1991. To test whether this kind of pragmatic effect could explain the observed difference in response times we collected judgments of felicity for the statement used in our study.

We collected judgments of the felicity of each of the belief and knowledge ascriptions in each of the three belief conditions for all twelve of the scenarios used in the previous experiments. This data allow us to ask both whether we find the predicted pragmatic difference in cases of true belief and whether the response time effect in the other two cases can also be accounted for by differences in felicity.

Methods

Participants. 537 participants ($M_{age} = 33.97$, $SD_{age} = 10.95$; 250 females) were recruited through Amazon Me-

chanical Turk.

Stimuli and procedure The procedure used was similar to that of Experiment 1a-b except that participants were asked to make a judgment about whether the belief or knowledge ascription ‘sounded weird’ or ‘sounded normal’ on a Likert scale (in addition to indicating whether the statement was true or false). Participants were also given unlimited time to respond. Prior to completing these trials, participants were trained on the task and completed four practice trials using statements that were clearly felicitous or clearly infelicitous.

Results and discussion

Felicity judgments for trials on which participants correctly assessed the truth of the knowledge and belief statements were analyzed as in the previous studies. This analysis revealed a main effect of Belief Condition, $\chi^2(2) = 25.588$, $p < .001$, a main effect Ascription Type, $\chi^2(1) = 42.238$, $p < .001$, and a significant interaction effect, $\chi^2(2) = 28.257$, $p < .001$. As predicted by theories of pragmatic implicatures, planned comparisons revealed a large difference in felicity judgments in the True Belief condition, $t(247) = 7.20$, $p < .001$, $d = 0.629$. While not theoretically predicted, we also observed a small but significant difference in the Ignorance condition, $t(228) = 3.29$, $p = .001$, $d = 0.209$. We did not observe a significant difference in the Different Belief condition, $t(231) = -1.23$, $p = .219$, $d = 0.080$.

Accordingly, we next asked whether these differences in felicity could account for the difference in response times, starting with the True Belief condition, where this kind of pragmatic difference would be expected. We first computed the mean felicity rating for each of the 24 knowledge and belief ascriptions and included these in the linear mixed-effects model (this is equivalent to ‘controlling’ for differences in felicity of each statement). We then asked whether the effect of Ascription Type on response times persisted after first modeling differences in felicity, and found that it did not, $\chi^2(1) = 0.942$, $p = .332$. In short, this result suggests that felicity judgments roughly tracked the differences in response times between knowledge and belief ascriptions, and may account for them.

Given this pattern, we next asked whether felicity judgments could also account for the difference in response times in the Ignorance and Different Belief scenarios, where it would not theoretically be predicted. We performed a similar analysis, but found that the effect of Ascription Type on response times persisted, $\chi^2(1) = 18.359$, $p < .001$. This result suggests that differences in the felicity of the knowledge vs. belief ascriptions in these cases does not explain the difference in response times. Using this same analysis approach, we tested whether the Belief Condition \times Ascription Type interaction effect remained. It did, and therefore is not accounted for by differences in felicity, $\chi^2(1) = 8.446$, $p = .004$.

We now take up the matter of why this interaction occurs—specifically, why we do not observe a difference in response times to belief versus knowledge in the different belief condi-

tion (but do in the other conditions). Possibly, comprehending the vignettes in which an agent has a false belief requires one to specifically understand the agent’s mind in terms of what the agent believes. In other words, the stimulus may bias participants to preferentially compute a belief state rather than a knowledge state. If correct, this would then specifically facilitate fast truth-value assessments of belief ascriptions but not knowledge ascriptions. In line with this explanation, we found that participants’ evaluations of belief ascriptions were in fact faster in the Different Belief condition ($M = 2552.53$, $SD = 634.01$) than in the Ignorance condition ($M = 2641.09$, $SD = 664.85$), $t(405) = -2.379$, $p = 0.0178$. By contrast, participants’ evaluations of knowledge ascriptions were actually slower in the Different Belief condition ($M = 2537.75$, $SD = 647.88$) than in the Ignorance condition ($M = 2461.32$, $SD = 645.40$), $t(478) = 2.549$, $p = 0.011$.

Together, Experiments 1 and 2 provide evidence that people are able to evaluate what others know before they are able to evaluate what they believe, and that this difference cannot be fully accounted for by pragmatic differences. Moreover, there is some evidence that this difference is attenuated when people are required to represent an agent’s beliefs in order to understand the stimulus (e.g., because it involves a false belief representation).

A general limitation of the previous studies is that they were conducted solely in English, and thus it is possible that the observed differences arise from some idiosyncratic feature of these English terms rather than reflecting the underlying cognition involved in representing agents as knowing or believing some proposition. For example, in English the term ‘know that’ is ≈ 1.52 times more frequent than ‘think that’. We take up this issue in the following experiment by asking whether the pattern we observed in Experiment 1 occurs cross-linguistically.

Experiment 3

We next conducted a similar experiment in French using ‘savoir’ instead of ‘know’, and ‘penser’ instead of ‘think’. French provides a particularly strong test case because, unlike English, the French term ‘penser que’ is ≈ 1.49 times more frequent than ‘savoir que’,⁴ and thus faster evaluations of knowledge ascriptions in French could not be explained by lexical frequency.

Methods

Participants 150 participants ($M_{age} = 37.70$, $SD_{age} = 12.16$; 83 females) were recruited and paid through Foule Factory (<https://www.foulefactory.com/>).

Stimuli and procedure

The methods and procedures in this experiment were identical to that of Experiment 1a-b, except that the study was translated into French, and the English names were replaced

⁴Lexical frequency was computed using Google Ngram for the most recent year available at the time of calculation (2007-2008).

with more typical French names. Thus, for example, instead of evaluating the truth or falsity of (1) or (2), participants evaluated the truth or falsity of (3) or (4).

- (3) Nora sait qu'elle regarde Neptune.
- (4) Nora pense qu'elle regarde Neptune.

Results and discussion

As in Experiment 1, data were excluded at the participant- and trial-level, and then analyzed using linear mixed-effects models. This revealed a marginal effect of Belief Condition, $\chi^2(2) = 5.513, p = .064$, and a highly significant main effect of Ascription Type, $\chi^2(1) = 25.351, p < .001$, such that participants were faster to correctly assess the truth of statements about what the agent knows ($M = 2565.27, SD = 722.63$) than statements about what the agent thinks ($M = 2729.57, SD = 719.52$). Additionally, there was again a significant Belief Condition \times Ascription Type interaction, $\chi^2(2) = 6.587, p = .037$ (see Figure 2).

Collectively, the previous experiments present cross-linguistic evidence that human adults are able to evaluate the truth of knowledge ascriptions prior to equivalent belief ascriptions. They also provide evidence against explaining this difference in terms of word frequencies or pragmatics, and instead suggest that the difference may reflect some aspect of underlying theory of mind processing. A final untested question is which aspect of knowledge and belief representation results in the observed difference in response times. We test one possibility in a final experiment.




Figure 2: Mean response time for evaluations of knowledge ascriptions (dark bars) and belief ascriptions (light bars) as a function of Belief Condition. Error bars depict ± 1 SEM.

Experiment 4

As discussed in the introduction, one fundamental difference between knowledge and belief is that knowledge is factive but belief is not. To ask whether this difference may underwrite the observed difference in response times between knowledge and belief, we next investigated response times for an independent set of factive and non-factive mental state ascriptions.

Methods

Participants 250 participants ($M_{age} = 33.432, SD_{age} = 9.32$; 126 females) were recruited and paid through Amazon Mechanical Turk.

Stimuli and procedure

The methods and procedures in this experiment were similar to the preceding studies except that the term ‘know’ in the mental state ascription was replaced by one of a set of factive attitude verbs (‘realize’, ‘is aware’, ‘understand’, ‘recognize’), and the term ‘think’ was replaced by a set of non-factive attitude verbs (‘believe’, ‘guess’, ‘assume’, ‘imagine’). Thus, for example, instead of evaluating the truth or falsity of (1) or (2), participants may have evaluated the truth or falsity of (5) or (6), respectively. Critically, these factive and non-factive terms were chosen such that the non-factive terms were both shorter in length and more frequent in use than the factive terms, such as:

- (5) **Factive:** Mira recognizes that she is looking at Neptune.
- (6) **Non-factive:** Mira believes that she is looking at Neptune.

Results and discussion

As in the previous experiments, data were excluded at the participant- and trial-level, and then analyzed using an identical set of linear mixed-effects models. This revealed no main effect of Belief Condition, $\chi^2(2) = 4.014, p = .134$ and no Belief Condition \times Ascription Type interaction, $\chi^2(2) = 0.955, p = .620$. We did, however, observe a highly significant main effect of Ascription Type, $\chi^2(1) = 11.127, p < .001$, such that participants were faster to correctly assess the truth of ascriptions involving factive attitudes ($M = 2362.83, SD = 610.68$) than ascriptions involving non-factive attitudes ($M = 2433.10, SD = 640.03$) (see Figure 3). While still tentative, these results provide some initial evidence that the response-time difference observed for evaluations of knowledge and belief ascriptions may generalize to a broader set of factive and non-factive mental state ascriptions.

General Discussion

The experiments we report seek to clarify the relationship between evaluations of knowledge and belief in human adults. To recap, Study 1 found that people can accurately evaluate others’ knowledge before they can accurately evaluate their beliefs. Study 2 demonstrated that this pattern cannot be not explained by pragmatic differences. Study 3 found that this pattern occurs cross-linguistically and is unlikely to be accounted for by differences in word frequency. Finally, Study 4 provided evidence that this difference in response times generalized to the larger class of factive and non-factive attitudes. Together, these experiments provide new evidence that correct evaluations of knowledge may occur faster to than equivalent evaluations of beliefs, and thus that knowledge assessment may often not depend on prior evaluations of belief.




Figure 3: Mean response time for evaluations of factive (dark bars) and non-factive (light bars) mental state ascriptions as a function of Belief Condition. Error bars depict ± 1 SEM.

Our findings fit well with the growing evidence from a number of different fields that representations of knowledge and other factive attitudes are more basic than than representations of non-factive attitudes like belief (Nagel, 2017; Williamson, 2002; Phillips & Norby, 2018). The capacity for factive mental state representation emerges earlier at a phylogenetic level (Martin & Santos, 2014), an ontogenic level (Luo & Johnson, 2009), a linguistic level (Bretherton & Beeghly, 1982), and also at the level of online processing.

Acknowledgments

We would like to thank Natalie Collins for her assistance with research. This research was supported by Grant N00014-14-1-0800 from the Office of Naval Research to JP and FC.

References

Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a theory of mind? *Cognition*, *21*(1), 37–46.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.

Bates, D., Maechler, M., Bolker, B., Walker, S., et al. (2014). lme4: Linear mixed-effects models using eigen and s4. *R package version*, *1*(7).

Bretherton, I., & Beeghly, M. (1982). Talking about internal states: The acquisition of an explicit theory of mind. *Developmental Psychology*, *18*(6), 906.

Call, J., & Tomasello, M. (2008). Does the chimpanzee have a theory of mind? 30 years later. *Trends in cognitive sciences*, *12*(5), 187–192.

Dennett, D. (1987). *The intentional stance*. Cambridge: MIT Press.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, *23*(6), 121–123.

Heim, I. (1991). Artikel und definitheit. *Semantik: Ein Internationales Handbuch der Zeitgenössischen forschung*, 487–535.

Ichikawa, J. J., & Steup, M. (2016). The analysis of knowledge. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Winter 2016 ed.). Metaphysics Research Lab, Stanford University.

Kiparsky, P., & Kiparsky, C. (1970). Fact. In M. Bierwisch & K. E. Heidolph (Eds.), *Progress in linguistics: A collection of papers*. Walter de Gruyter GmbH & Co. KG.

Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, *354*(6308), 110–114.

Lehrer, K. (1968). Belief and knowledge. *The Philosophical Review*, 491–499.

Leslie, A. M. (1987). Pretense and representation: The origins of "theory of mind." *Psychological review*, *94*(4), 412.

Luo, Y., & Johnson, S. C. (2009). Recognizing the role of perception in action at 6 months. *Developmental Science*, *12*(1), 142–149.

Mar, R. A., Tackett, J. L., & Moore, C. (2010). Exposure to media and theory-of-mind development in preschoolers. *Cognitive Development*, *25*(1), 69–78.

Martin, A., & Santos, L. R. (2014). The origins of belief representation: Monkeys fail to automatically represent others beliefs. *Cognition*, *130*(3), 300–308.

Martin, A., & Santos, L. R. (2016). What cognitive representations support primate theory of mind? *Trends in Cognitive Sciences*, *20*(5), 375–382.

Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that p without believing that p. *Noûs*, *47*(2), 371–384.

Nagel, J. (2017). Factive and nonfactive mental state attribution. *Mind & Language*, *32*(5), 525–544.

Onishi, K. H., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, *308*(5719), 255–258.

Phillips, J., & Norby, A. (2018). Factive theory of mind. (under review)

Radford, C. (1966). Knowledge by examples. *Analysis*, *27*(1), 1–11.

Santos, L. R., Nissen, A. G., & Ferrugia, J. A. (2006). Rhesus monkeys, *Macaca mulatta*, know what others can and cannot hear. *Animal Behaviour*, *71*, 1175–81.

Snyder, C. R., Harris, C., Anderson, J. R., Holleran, S. A., Irving, L. M., Sigmon, S. T., ... Harney, P. (1991). The will and the ways: Development and validation of an individual-differences measure of hope. *Journal of personality and social psychology*, *60*(4), 570.

Vouloumanos, A., Martin, A., & Onishi, K. H. (2014). Do 6-month-olds understand that speech can communicate? *Developmental Science*, *17*(6), 872–879.

Williamson, T. (2002). *Knowledge and its limits*. Oxford University Press.