Veridicality and the acquisition of *think*

Peter van Elswyk, Northwestern University

Forthcoming in *Linguistics and Philosophy*; Please cite the official version

Across numerous languages, the attitude verb *think* is learned later than other attitude verbs like *want*. But why? This essays advances a new hypothesis: children initially treat *think* as a veridical yet non-factive verb akin to a class of verbs I call confirmatives. This hypothesis is argued to better explain existing data that troubles other hypotheses, and to find support from the ease with which children represent knowledge but not belief.

1 The acquisition of *think*

How children learn natural language is a puzzle. From clues found in the context, children have to pair a word with a semantic meaning. But the context almost always does not give enough clues (Quine, 1960). As a result, the available clues are compatible with a heap of word-meaning pairs with only one or a few being correct. And yet, children eventually succeed in finding the needle in the haystack. They find the correct pairing(s).

How is this accomplished? The first part of an answer is that context involves more than the physical environment. It also includes the linguistic context (Carey, 1978; Gleitman, 1990). For a word w, the linguistic context divides into four sub-contexts:

LEXICAL CONTEXT
The words with which w co-occurs.

SYNTACTIC CONTEXT
The syntactic features with which w co-occurs.

PRAGMATIC CONTEXT
The pragmatic uses of language with which w co-occurs.

DISCOURSE CONTEXT
The discourse segment with which w co-occurs.

In principle, children can find clues in all of these varying sub-contexts of the linguistic context.

The second part of an answer is social cognition. Just as we have psychological abilities for navigating a complex physical environment, we also have abilities for

navigating a complex social environment too. Such abilities include joint attention and mindreading, *i.e.*, the ability to track the mental states of another that are causally relevant to their behavior. Though not yet fully developed, children use these abilities to closely monitor speakers while learning language (Bloom, 2002; Nadig and Sedivy, 2002; Khu et al., 2020). The question of a word's acquisition therefore boils down to how social cognition and clues from the linguistic context jointly enable the child to learn the word.

This paper's focus is how children learn the attitude verb *think*, especially in contrast to *want*. Both verbs denote mental states that play a starring role in our mental lives. However, the verbs are learned in stark asymmetry. Across numerous languages, children have an adult-like mastery of *want* at three years, but lack similar mastery of *think* until after they turn four (Wellman and Woolley, 1990; Perner, 1991; Tardif and Wellman, 2000; Perner et al., 2003; Harrigan et al., 2018). The difference is noticeable in how children respond to attitudes with false complements. Suppose it is clear that Pip is playing with the dog. Then both of these attitude reports are incongruent with reality.

(1) Margaret wants Pip to take a bath.

(2) Margaret thinks that Pip is taking a bath.

In (1), Margaret wants what is not the case. In (2), Margaret has a false belief about what is the case. Children struggle with verbal false belief tasks until ages four to five (Wimmer and Perner, 1983; Wellman et al., 2001). Even though the *think*-attribution is true, they reject sentences like (2) when the complement is false. However, they do not struggle with reality-incongruent *want*-attributions. Sentences like (1) are consistently accepted around age three.

A standard theoretical commitment is that the acquisition of *think* orients around a mistake. Initially, a child makes a mistake about how to understand *think*-attributions. This is why its acquisition is delayed in contrast to mental state verbs like *want*. For example, some hypothesize that children do not initially understand the syntax of complementation (Diessel and Tomasello, 2001; de Villiers and Pyers, 2002), others suggest children mislearn *think* as a factive verb like *know* (Johnson and Maratsos, 1977; Abbeduto and Rosenberg, 1985), and some propose that children go into "pragmatic overdrive" to misinterpret how the *think*-attribution is being used by the speaker (Hacquard and Lidz, 2019, 2022). This paper will advance a new hypothesis that similarly identifies a mistake.

In pursing a mistake-oriented hypothesis, the question of *think*'s acquisition divides into two distinct questions. The first question concerns what mistake is consistently made. The second question concerns how the mistake is consistently corrected within the relevant timeframe. Notably, an answer to the second question is often overlooked. The general nature of the correction is made clear by the hypothesis—learning complementation, untangling *think* from *know*, not going into pragmatic overdrive—but how children manage to achieve this correction in the relevant timeframe is left undiscussed.

Making progress on the acquisition of *think* requires us to answer both questions. We need to identify what leads children away from correctly understanding *think*-attributions and what pushes them in the right direction around their fourth birthday. Even still, investigating these questions separately is a worthwhile endeavor given the various ways answers can be combined. Though answers to the first question will place some constraints on how the second question is answered, it is plausible that an answer to the second question is compatible with competing answers to the first question. In other words, an explanation of how children are pushed in the right direction may be compatible with different explanations of how they were lead astray in the first place.

With these questions distinguished, the aim of this paper is modest. It addresses only the first question by advancing the VERIDICALITY HYPOTHESIS, a view according to which children initially mistake *think* for a verb that is veridical and yet non-factive. In the next section, I present the positive case (§2). Then I compare it to THE PRAGMATIC SYNTACTIC BOOTSTRAPPING HYPOTHESIS, advanced by Hacquard and Lidz (2019, 2022) and favored by Dudley et al. (2015), Lewis et al. (2017), and Harrigan et al. (2019) (§3). Both hypotheses only address the first question. However, the latter hypothesis arguably provides the leading answer to the first question. So showing that the veridicality hypothesis outperforms it offers another way to motivate that the hypothesis is a live option.

My approach to motivating the veridicality hypothesis is modest as well. I intend to show that the hypothesis is well-supported by existing empirical data, as opposed to earning plausibility through new findings. Accordingly, I will provide a reassessment of existing data as opposed to introducing data uncovered through new experiments. Since previous work has overlooked the hypothesis, it is important to see what is clarified by the data we already have. Nevertheless, I will end the paper by discussing some of the paths for future experimentation that the veridicality hypothesis paves (§4).

2 The veridicality hypothesis

2.1 Veridicality as a property of verbs

A verb $\mathcal{V}$ is veridical if and only if $\mathcal{V}(p)$ entails p. Veridical verbs are closely related to factive verbs (Kartunnen, 1971; Egré, 2008; Spector and Egré, 2015). All factive verbs are veridical but not all veridical verbs are factive. Factive verbs differ by backgrounding the veridicality entailment as a presupposition. As a result, the veridicality entailment for factives projects past the scope of entailment-canceling operators (Chierchia and McConnell-Ginet, 2000). To illustrate, compare *prove* with *know* below.

(3) (a) {I / the professor} knew that $\sqrt{400}$ is 20.
(b) {I / the professor} proved that $\sqrt{400}$ is 20.

(4) (a) {I / the professor} didn't know that $\sqrt{400}$ is 20.
(b) {I / the professor} didn't prove that $\sqrt{400}$ is 20.

While *know* is a factive verb, *prove* is just a veridical verb. As a result, the two verbs both intuitively entail their complement clause in (3a) and (3b), or that $\sqrt{400}$ is 20. But the two verbs can be distinguished by the entailments they license under negation. Negation blocks mere entailments; it doesn't block presuppositions. Accordingly, *know* still licenses the entailment in (4a). In contrast, the veridicality entailment licensed by *prove* disappears.

Veridical but non-factive verbs include *prove*, *confirm*, and *show*.[1] These form a natural class of verbs indicating that the grammatical subject performed, performs, or will perform an action to reveal the truth of proposition contributed by the verb's complement (Zuchewicz, 2020). I will call them CONFIRMATIVES. Though these verbs are often used to report the confirming act of another (e.g. *Margaret proved that*...), the rough-and-ready notion of revealing here does not require the confirming act to

[1] Anand and Hacquard (2014) argue that *confirm* is not veridical because the phrase *falsely confirm* is well-attested. However, there is good reason to doubt this diagnostic. First, *falsely* $\mathcal{V}$ where $\mathcal{V}$ is either a speech act verb or a confirmative is frequently used in legal contexts as a term of art to indicate that $\mathcal{V}$-ing occurred fraudulently. A detailed corpus study is required to substantiate whether most uses are legal uses. But this appears to be the case for confirmatives. Second, *falsely* can appear with common factive verbs as well. For example, *falsely realize*, *falsely remember*(*ed*), and even *falsely know* are all attested. Here is an example of the latter: *The vast majority of people in the US know next to nothing about the Episcopal Church, or falsely know it as the church that Henry the VIII started so that he could divorce as many wives as he wanted* (https://saintthomasepiscopal.org/sermon-for-clergy-day-june-2016/). As such, appearing alongside *falsely* is not a reliable diagnostic for non-veridicality. In such uses, the meaning of *falsely* is either non-standard legalese, or best interpreted as akin to *mistakenly takes oneself to* $\mathcal{V}$.

be public to someone other than the subject. A person can use a sentence like (3b) to report what they revealed in private.

Confirmatives are not speech act verbs. As just noted, they can be used to denote confirmations that happened in private with nobody around to be an addressee. They can also be used to denote non-speech actions. For example, suppose a professor demonstrates in total silence how to figure out a square root by first simplifying the square root into perfect square factors such as how $\sqrt{400}$ is $\sqrt{24 \times 16}$. Their demonstration could be reported with a confirmative (e.g. *The professor showed that*...). However, confirmatives can be used to report speech acts. Speech acts like assertion do meet our rought-and-ready notion of an act that reveals a proposition's truth. For example, suppose we've heard that $\sqrt{400}$ is 20. But we are still curious whether that's true. We want verification. Were the previously imagined professor to assert that the answer is 20, his assertion could be felicitously reported with *The professor confirmed that* $\sqrt{400}$ *is 20*.

Confirmatives are not attitude verbs either. Their core semantic meaning is to denote a truth-revealing action as opposed to a mental state. However, a closer look reveals that confirmatives encode speaker belief like factive verbs do. A standard diagnostic for whether an expression $\mathcal{V}(p)$ entails belief in p is to consider the felicity of the conjunction $\ulcorner \mathcal{V}(p)$, but I don't believe $p \urcorner$ (Faller, 2002; Papafragou, 2006; Murray, 2017). If the conjunction is infelicitous, this is evidence that $\mathcal{V}(p)$ licenses the inference that the speaker believes p. That doxastic inference is what causes the infelicity. However, if the conjunction is felicitous, this is evidence that $\mathcal{V}$ licenses no such inference. With that in mind, compare *heard*, *realize*, and *prove* in the following conjunctions.

(5) {I / The professor} heard that $\sqrt{400}$ is 20. But I don't believe that.

(6) # {I / The professor} realized that $\sqrt{400}$ is 20. But I don't believe that.

(7) # {I / The professor} {proved / confirmed / showed} that $\sqrt{400}$ is 20. But I don't believe that.

Verbs reporting what was heard are fully felicitous with subsequent disavowals of belief (AnderBois, 2004). Factive verbs are not. Of interest is that *prove* patterns with *realize* as opposed to *heard*. It is infelicitous in conjunctions of the form $\ulcorner \mathcal{V}(p)$, but I don't believe $p \urcorner$. The same goes for the other confirmatives. Such infelicity shows that confirmatives entails that the speaker believes p. This is an intuitive result given

our gloss of confirmatives. If a speaker performs an action to reveal p's truth, they give themselves a reason to believe p.[2]

Importantly, the doxastic inference of confirmatives is not presupposed. To see as much, it is instructive to compare confirmatives with the verb phrase *be right*. It presupposes the doxastic entailment but not the veridicality entailment. Compare the following.

(8) Is Margaret right that Pip took a bath?

(9) Did the professor {prove / show / confirm} that Tait's conjecture is true?

A question operator is another entailment-canceling operator that blocks mere entailments but which presuppositions project past. In (8), the doxastic inference projects. That Margaret believes that Pip took a bath is not in question. In contrast, the doxastic inference in (9) does not. Such a question can be felicitously asked when neither the professor nor the speaker believes that Tait's conjecture is true. For example, it can be asked in a context where the relevant parties are skeptical about the conjecture's truth. Confirmatives therefore constitute a unique semantic class of verbs. Lewis et al. (2017) claim that all verb phrases with the doxastic and veridicality entailments background at least one of the entailments. Confirmatives are an exception that presuppose neither.

2.2 The hypothesis

Let's now set the stage for the veridicality hypothesis. When children encounter attitude verbs like *think* and *want*, they engage in SYNTACTIC BOOTSTRAPPING (Gleitman, 1990; Gillette et al., 1999; Landau and Gleitman, 2009). These verb correlate differently with kinds of complements, number of arguments, and mood. The correlations are then used to discern the meaning of the verbs (Hacquard and Lidz, 2019). In English, for example, the verb *think* embeds a finite, declarative sentence whereas *want* does not. It embeds a non-finite complement akin to an

[2] A wrinkle is that the doxastic inference disappears in certain syntactic configurations. In particular, if a confirmative is combined with an inanimate but information-containing subject like *The book* and a prepositional phrase is added to the confirmative to indicate who observed the confirmation like *to Margaret*, then the inference is no longer licensed. For example, *The book proved to Margaret that* $\sqrt{400}$ *is 20, but I don't believe that* is a felicitous conjunction. Anand and Hacquard (2014) observe that the veridicality entailment also disappears in these configurations. As such, the doxastic inference appears to depend on the veridicality entailment for both factive verbs and veridical but non-factive verbs.

imperative sentence in at least some respects. Children learning other languages rely on different morphosyntactic clues to sort the verbs by complements. For example, Spanish speakers sort the verbs by the different moods of the complements, and Mandarin speakers rely on other morphosyntactic contrasts (Huang et al., 2022). As a result, *think* gets sorted into a semantic class of verbs that can be veridical or not whereas *want* gets sorted into a different class.

Dudley et al. (2017) undertook a corpus study to investigate how children encounter the verbs *think* and *know*. In the corpus investigated, children range from ages two to five with the average being 3.5. As such, the corpus provides a clear window into what clues children are encountering in the relevant age range. They found that children encounter *think* mostly with a first-person subject (67%), in the present tense (91%), and without negation (89%). Dudley et al. (2017, 613) conclude that *think* is mostly used to discuss the speaker's own attitudes and that "there may be few instances where a child could observe that *think* can be used to describe false [attitudes] and is thus non-veridical." Accordingly, children are overwhelmingly seeing *think* as having the doxastic entailment, and, given few clues from the lexical or syntactic sub-contexts, that it lacks the veridicality entailment. In this situation, children have to sort *think* into a class of verbs that is doxastic and veridical or a class that is doxastic and not veridical.

According to the veridicality hypothesis, children sort *think* into the veridical class. As a result, children mistakenly learn *think* as a verb that has the core semantic properties of confirmatives.[3] The process of correctly acquiring the meaning of *think* is therefore the process of unlearning the veridicality of *think* to be left with a verb that is merely doxastic. This learning process requires navigating the asymmetric relationship between veridicality and doxasticity as components of a verb's meaning. Though all veridical verbs are doxastic, not all doxastic verbs are veridical. That latter semantic class is where *think* belongs.

Failing verbal false belief tasks inevitably results until veridicality is unlearned. When children encounter a true *think*-attribution with a false complement, the attribution is rejected because the attribution is treated as having a veridicality entailment. That the doxastic component is true is not sufficient for the truth of

[3] Or, at least, the hypothesis is that *think* has the core semantic properties of confirmatives in the past tense. Confirmatives have an incrementality to their meaning in other tenses. For example, *The professor is proving that* $\sqrt{400}$ *is 20* does not have the same ring of veridicality. Veridicality seems to be associated with the successful completion of the proof. It is the proof not the proving that secures truth. See Zuchewicz (2020) for a detailed discussion of the tense and aspect of confirmatives in Polish.

the attitude report. The same issue does not arise with *want* because it cannot be veridical. Veridicality is exclusively a semantic property of verbs with declarative and/or interrogative complements.

The veridicality hypothesis therefore offers a clear explanation for why *think* is learned late, especially in contrast to verbs like *want*. That explanation is appreciably simple by isolating one semantic feature that children mislearn. But it raises a question: why do children learn *think* as having the core semantic properties of confirmatives?

2.3 Veridicality from congruent mindreading

In the absence of linguistic clues that the verb is veridical or not, children have to base their sort of *think* on something else. This is where social cognition kicks in.[4] It compels children to treat *think* as veridical.

The ability to represent another's beliefs is one form of mindreading. Another form is the ability to represent another's knowledge. Representing knowledge has been called FACTIVE MINDREADING and contrasted with NON-FACTIVE MINDREADING where one mindreads another as merely believing (Phillips and Norby, 2021; Westra and Nagel, 2021). In the present context where we are distinguishing mere veridicality from factivity, this name is misleading. What these authors intend to draw attention to is that attributing knowledge includes a commitment to the the truth of the complement. In contrast to representing someone as believing, to represent someone as knowing a proposition is to represent them as being connected to the truth. The factive/non-factive distinction for mindreading is itself an umbrella term for the different names that psychologists give to the representation of knowledge during mindreading. Alternatives to *factive* include *reality-congruent*, *veridical*, *awareness*, *visual mindreading*, and *level-1 perspective-taking*. To avoid confusion, I opt to use CONGRUENT/NON-CONGRUENT MINDREADING.

An impressive body of evidence reveals that knowledge representation is prior to belief representation (Phillips et al., 2021). One way to illustrate the priority with respect to development is to consider when children can succeed on nonverbal mindreading tasks. Around 18-24 months, children are able to engage in congruent mindreading. This has been shown with success on violation-of-expectation tasks

[4] Others have similarly suggested that limitation in a child's social cognition are to blame. For example, see Perner (1991), Tardif and Wellman (2000), Leslie et al. (2004), and Steglich-Petersen and Michael (2015). But none of these authors chalk up the difference to congruent versus non-congruent mindreading.

(Träuble et al., 2010; Yott and Poulin-Dubois, 2016), action-anticipation tasks (Surian et al., 2007; Sodian et al., 2007; Luo and Baillargeon, 2007; Luo and Johnson, 2009), and active-helping tasks (Hamlin et al., 2013). A similar cluster of non-verbal tasks was one thought to converge on the conclusion that children could represent false belief in the same developmental window (Onishi and Baillargeon, 2005; Southgate et al., 2007; Buttelmann et al., 2009). However, as Poulin-Dubois et al. (2018) detail, the relevant studies have failed to replicate by numerous labs outside the original ones. Altogether, the conclusion that is increasingly drawn is that children have adult-like knowledge representation but cannot yet represent belief. For example, Powell et al. (2018, 40), after failing to replicate nonverbal false belief tests, suggest the results cumulatively may "reflect the veridical abilities of 18-month-old infants, who may track others' knowledge and ignorance but may not consistently represent the contents of others' beliefs."

Another way to illustrate the developmental priority of congruent over non-congruent mindreading is to consider success on the Theory-of-Mind scale developed by Wellman and Liu (2004). This scale consists in a series of verbal tasks aimed to evaluate a child's ability with a particular form of mindreading including, among a few others, knowledge representation and false belief representation. Though there is cultural variation across children from the United States, Australia, China, Iran, and Turkey in what exact order children proceed through these tests, success on knowledge tasks always precedes success on false belief tasks (Shahaeian et al., 2011; Doenyas et al., 2018; Ilgaz et al., 2022).[5]

Knowledge representation is the default form of mindreading that is deployed in conversation (Nagel, 2013; Westra and Nagel, 2021). When a speaker uses a declarative to perform an assertion, the hearer represents them as knowing what was asserted. Likewise, when a speaker asks an information-seeking question, the hearer represents them as not knowing the answer or answers to the question (Brown-Schmidt and Fraundorf, 2015; Aguirre et al., 2022). My proposal is that this default is also operative when children are attempting to learn attitude verbs with declarative complements.

When children encounter *think* used primarily in the first-person to report an

[5] Both of these illustrations have limitations. The first turns on the replicable success on various nonverbal tasks for congruent but not for non-congruent mindreading. The current state of replication may turn out to not be a lasting set of results, especially if there is innovation in nonverbal tasks. The second illustration turned on the Theory-of-Mind scale. This scale is controversial in part because it introduces confounds related to the pragmatic abilities of children. See Westra and Carruthers (2017).

attitude, they misrepresent that attitude as knowledge. To attribute knowledge of *p* is to regard *p* as true and believed by whoever knowledge is attributed to. But these are the core semantic properties of confirmatives (§2.1). In defaulting to knowledge representation while attempting to learn *think*, children therefore treat the attitude verb as being akin to a confirmative.

2.4 The hypothesis's predictions

The veridicality hypothesis makes specific predictions about how children will treat *think*. We can therefore evaluate the hypothesis by considering how previous research speaks to these predictions. Let's start with the combination of truth-values for the matrix and complement clauses. On the hypothesis that the verb is veridical, young children will reject a *think*-attribution when the complement clause is false. On the hypothesis that the verb is doxastic, children will also reject a *think*-attribution when the matrix clause is false. Both need to be true to be accepted. This yields four predictions:

	MATRIX CLAUSE *I think*	COMPLEMENT CLAUSE *that p*	VERIDICALITY HYPOTHESIS PREDICTION
(1)	True	True	Accept
(2)	True	False	Reject
(3)	False	True	Reject
(4)	False	False	Reject

Figure 1

Predictions (1) and (4) are the uninteresting predictions because they are well-confirmed, and do not distinguish the veridicality hypothesis from rival ones. The interesting predictions are the middle two.

The difficulty with verbal false belief tasks previously discussed shows that prediction (2) is accurate (§1). This is the prediction the hypothesis is engineered to explain. Prediction (3) is what distinguishes the veridicality hypothesis from the COMPLEMENTATION HYPOTHESIS according to which children acquire *think* late because they struggle with the syntax of complementation (Diessel and Tomasello, 2001; de Villiers and Pyers, 2002). If children evaluate *think*-attributions only according to the truth of the complement clause, for example, then children will mostly be indifferent to the falsity of the matrix clause. They will accept the attribution when the complement is true. The veridicality hypothesis predicts no such indifference.

Since *think* still has a doxastic component, false attributions will be rejected by children even with true complements.

Previous research confirms that the veridicality hypotheses gets this prediction right. Lewis et al. (2017) investigated how children as young as three react to *think*-attributions in this condition. These attributions were entertained in contexts where the matrix clause alone was much more relevant than the complement clause in that it answers the question under discussion in the discourse. Setting the context in this way offered a way to control for children giving their attention to the complement clause by default. In this context, they found that children reliably rejected attributions with true complements but false matrix clauses. Accordingly, children are not indifferent to the matrix clause. They are sensitive to whether the subject believes the complement.

In making these predictions, the veridicality hypothesis is similar to an earlier hypothesis that *think* was initially mislearned as a factive verb like *know* or *realize* (Johnson and Maratsos, 1977; Abbeduto and Rosenberg, 1985). Let's call this the FACTIVITY HYPOTHESIS. The difference between these two hypothesis is subtle. Only factives presuppose the entailment. As a result, the entailment projects past the scope of entailment-canceling operators like negation only for factives. This is testable and has been tested indirectly. Dudley et al. (2015) explored whether three-year-olds infer p from ⌜Lambchop doesn't think p⌝ versus ⌜Lambchop doesn't know p⌝. Since p projects past negation, children should show a similar tendency to infer p from both constructions. But they do not show a similar tendency. Instead, all of the children in their study appear to understand that *think* is non-factive with some grasping that *know* is factive.

At the outset, I noted that a mistake-oriented explanation of *think*'s acquisition answers two questions. The first question is about what mistake is consistently made, and the second is about how the child consistently stops making that mistake. The veridicality hypothesis only answers the first question in detail. However, it does carry partial commitments for how to answer the second question. If difficulty with non-congruent mindreading contributes to mislearning *think*, children need to become better at non-congruent mindreading by age four. It predicts they will, and we can evaluate this prediction.

The relevant development evidence suggests this prediction is vindicated. By age four, children appear to have this ability, as illustrated by success on nonverbal false belief tasks. These tasks include change-of-location tasks applied to children at five years of age and primates (Call and Tomasello, 1999; Krachun et al., 2009). They also

include the action-anticipation tests mentioned earlier. For example, Kammermeier and Paulus (2018) attempted to replicate the finding that children as you as young as three could pass such nonverbal tasks. That replication failed, which is line with the developmental priority of congruent mindreading. However, they did find that children increasingly performed better as they approached, and then passed, their fourth birthday.

3 The pragmatic syntactic bootstrapping hypothesis

3.1 An alternative hypothesis

The leading explanation of why *think* is learned late is what Hacquard and Lidz (2019, 2022) call the PRAGMATIC SYNTACTIC BOOTSTRAPPING HYPOTHESIS, or what I will just call the BOOTSTRAPPING HYPOTHESIS. To complete my case for the veridicality hypothesis, I will now discuss how it compares to the bootstrapping hypothesis. The veridicality and bootstrapping hypothesis begin in the same place: with children sorting *think* and *want* into different semantic classes based on the syntactic complements that each attitude embeds.[6] From here, the hypotheses head in different directions. Where the veridicality hypothesis posits that children make a semantic mistake in mislearning *think*, the bootstrapping hypothesis posits that children make a pragmatic one in misunderstanding its uses.

Most children encounter *think* in the present tense with a first-person subject (Bloom et al., 1989; Dudley et al., 2017). In such a configuration, verbs like *think* are frequently used to hedge assertions. An example is the reply in (12b). The question under discussion does not concern what the speaker is thinking. It targets who is taking a bath.

(10) (a) Who is taking a bath?
(b) I think Pip is taking a bath.

However, (12b) is not therefore irrelevant. It is indirect: the complement answers the question and the *I think* hedges how this answer is proffered by the speaker. However, not all uses of *think* are indirect assertions. Consider the same reply with another question.

[6] This is by design. With respect to syntactic bootstrapping, the hypothesis best explains the data. As a result, the veridicality hypothesis is, in part, inspired by the bootstrapping hypothesis minus its pragmatic commitments.

(11) (a) Who do you think is taking a bath?
(b) I think Pip is taking a bath.

In this different context, the assertion is direct: *I think* contributes to an answer and the complement does not. According to the bootstrapping hypothesis, the mistake is that children treat sentences like (13b) as if they were (12b). In other words, they treat *think*-attributions as indirect, hedged assertions about p as opposed to direct, flat-out assertions that the speaker thinks p.

In treating *think*-attributions as indirect assertions, children are therefore compelled to reject the sentence when the complement is false even if the whole sentence is true. This is because the complement is AT-ISSUE or what has MAIN POINT STATUS (Urmson, 1952; Hooper, 1975; Simons, 2007). The *I think* is backgrounded such that ⌜I think p⌝ is treated as if it were just p. No similar mistake is made with *want*. Since it takes a non-finite complement akin to an imperative sentence that contributes a property, verbs like *want* can be used to perform indirect commands when in the present tense with a first-person subject. An example is *I want you to take a bath*. Here the addressee is indirectly directed to satisfy the speaker's desire by bringing about the property denoted by *to take a bath*. As a result, children cannot make the same mistake with *want* because *want*-attributions are never indirect assertions. Children are not thereby compelled to reject true *want*-attributions when the complement is reality-incongruent.

The bootstrapping hypothesis is similar to the complementation hypothesis. It identifies a fixation on complement clauses as the explanation for why children fail on verbal false belief tasks until age four. But it identifies a different cause for this fixation. It is not a difficulty with the syntax of complementation. Instead, it is a tendency to go into "pragmatic overdrive," to use Hacquard and Lidz's (2019) phrase, and interpret the complement as an indirect assertion even when it is not. But can children go into pragmatic overdrive? The next section provides reasons to be skeptical.

3.2 The demands of pragmatic overdrive

When attitude verb is used to hedge, the resulting speech act is different from a typical assertion in at least two ways: it is indirect and and it is weaker. It is indirect because the verb's complement is what's proffered as opposed to the proposition contributed by the whole sentence. It is weaker because the proposition is proffered with less force or oomph. When a speaker opts to use ⌜I think p⌝ over the unhedged p,

they are choosing to stand behind p less. We can therefore ask whether children are capable of understanding the indirection and weakness brought about by hedging. Let's start with the weakness effect.

An assertion is associated with the expectation that the speaker knows what they assert.[7] Hedging overrides this expectation by indicating that the speaker possesses a weaker attitude like mere belief (Benton and van Elswyk, 2020; van Elswyk, 2024). A construction like ⌜I think p⌝ is weaker, if used to hedged, than the mere assertion of p because it carries the scalar implicature that the speaker does not know p. Factive verbs cannot be used to hedge for the related reason that knowledge is the most general factive state (Williamson, 2000; Nagel, 2017). As a result, factive verbs do not license a scalar implicature. For illustration, compare (14b) and (15b). Only the first can be used to hedge.

(12) (a) Who is taking a bath?
(b) I think Pip is taking a bath. But I don't know that.

(13) (a) Who is taking a bath?
(b) # I realized Pip is taking a bath. But I don't know that.

The question of whether a child as young as three can understand the weakness effect of hedging can therefore be assimilated to the question of whether such children can understand scalar implicatures.

At present, the evidence is that they cannot. Though most research has focused on children's understanding of implicatures related to the scalar contrast between *some* and *all*, this research sets mastery no earlier than age four with age seven being a plausible benchmark too (Noveck, 2001; Papafragou and Musolino, 2003; Huang and Snedeker, 2009; Katsos et al., 2016). Such findings dovetail with results suggesting that children do not order *think* below *know* in strength until age four (Moore et al., 1989; Moore and Davidge, 1989; Falmagne et al., 1994; Naigles, 2000). Instead, the two attitude verb are treated as being indistinguishable in strength. Accordingly, it does not seem plausible that children can understand the weakness effect before they are four.

More generally, to understand the weakness effect requires children to understand that speakers are less reliable given that signaling weakened reliability is the

[7] This is the most commonly held view in philosophy but has less traction in linguistics. For discussion, see Williamson (2000), Benton (2011, 2016, forthcoming), Turri (2010, 2016), and van Elswyk and Benton (2023) for defense. See Turri (2017, 2018) for experimental research favoring the centrality of knowledge.

purpose of hedging. As van Elswyk (2023) argues, hedging enables speakers to cooperatively share unknown information by overtly marking their testimony as less reliable than an unqualified assertion. But children have a strong presumption that the speaker is reliable, especially if it is a parent (Corriveau et al., 2009). Zeroing in on three-year-olds, they are less selective than four-year-olds about who is reliable preferring all-or-nothing judgments of reliability to graded ones (Harris, 2007), they are less likely than four-year-olds to revise beliefs even when a trusted speaker is shown later to be unreliable (Scofield and Behrend, 2008), they will believe counterintuitive claims and are less likely than four-year-olds to voluntarily test those claims (Ronfard et al., 2017), and, even if they do test those claims and find them false, three-year-olds are less likely to modify their judgment of the speaker's reliability (Hermansen et al., 2021). As a result, understanding the weakness effect of hedging appears to require subtly in understanding the reliability of speakers that three-year-olds have not fully developed.

What about indirection? Here it depends on how children are determining that an instance of ⌜I think p⌝ is hedged. Hacquard and Lidz (2019, 92) appear to be committed to the view that children rely on intention-recognition, or as they put it, that "children's errors with *think* arise from children reading too much into the intentions of a speaker using a *think* sentence." However, intention-recognition is very cognitively demanding given that it involves recursive, higher-order mindreading (Kinderman et al., 1998; Launay et al., 2015; Oesch and Dunbar, 2017). The child needs to believe$_1$ that the speaker intends$_2$ them to know/believe$_3$ that the speaker is proffering the complement and not the matrix clause. This is minimally at least three orders of attitude representation. If representing the speaker as proffering involves attitude representation too—if to proffer p is to believe or know p, for example—then we are brought to four orders of attitude representation. But this is too demanding. For example, Liddle and Nettle (2006) found that ten and eleven-year-olds barely perform above chance on third-order mindreading tasks, and were at chance with fourth-order mindreading tasks. Children considerably younger than this will be worse off, especially given the difficulties with non-congruent mindreading noted earlier (§2.2).

The pragmatic overdrive posited by Hacquard and Lidz's bootstrapping hypothesis therefore appears to be beyond the ability of three-year-olds. Given known difficulty with scalar implicatures and their strong presumption that the speaker is reliable, young children do not appear to be able to understand the weakness effect generated by hedging uses of *think*. Children are also not yet able to engage in the

form of recursive mindreading required to identify via intention-recognition that a *think*-attributions is a hedged assertions.

The veridicality hypothesis does not share these challenges. Since it traces the late acquisition of *think* to a semantic mistake as opposed to a pragmatic one, it does not require that children have the pragmatic abilities at three that the bootstrapping hypothesis does. Instead, the veridicality hypothesis can explain some of the data that makes trouble for it. In particular, the hypothesis offers a natural explanation for why children treat *think* and *know* as being the same in strength until age four. They do so because they treat *think* as being akin to a confirmative. Since confirmatives and *know* both carry veridical and doxastic entailments, children cannot order them by what entails what. With respect to these two properties, they are mutually entailing. It is not until they unlearn that *think* is veridical that they can recognize that *know* is stronger.

4 Going forward

The veridicality hypothesis issues correct predictions about when children as young as three accept and reject *think*-attributions. The hypothesis also finds motivation in recent developmental psychology on congruent/non-congruent mindreading, and avoids problems had by the complementation, factivity, and bootstrapping hypotheses. As a result, it provides a plausible answer to the first question raised by a mistake-oriented explanation of acquisition. The mistake children make is learning *think* as a confirmative verb.

Where do we go from here? The next step is answering the second question, or how children correct the hypothesized mistake around their fourth birthday. The beginning of an answer was provided earlier (§2.4). The inability to engage in non-congruent mindreading around age three contributes to children mislearning *think*. However, that ability appears to be up-and-running around age four. But this fact does not provide us with a complete answer to the second question all by itself. Even if the cause of mislearning *think* is no longer present in the relevant timeframe, children still need to correct what they mislearned when the cause was present. Correcting their mistake requires shaving off the veridicality entailment from their understanding of the verb's meaning. What encourages them to do so is what an answer to the second question will clarify.

The second question has been beyond the scope of this paper. This means that the veridicality hypothesis does not provide a complete explanation of the acquisition

of *think*. But it is worth stressing that this is incompleteness is standard to mistake-oriented explanations of *think*'s acquisition. To the extent that this incompleteness is a liability for the veridicality hypothesis, it is a liability for other alternatives including, but not limited to, the bootstrapping hypothesis. For example, whatever is involved with going into pragmatic overdrive, children need to correct for that mistake around age four.

A plausible hypothesis does not merely outperform other hypotheses at explaining the phenomenon in question. It also opens up new routes by which to investigate the target phenomenon. To conclude, I will sketch some of the paths that the veridicality hypothesis encourages us to take.

One path forward is to investigate the acquisition of confirmative verbs such as *prove*, *confirm*, and *show*. The mindreading-based explanation for the veridicality hypothesis leads us to expect that at least the veridical and doxastic components of these verbs are not difficult to learn (§2.2). If three-year-olds are quick to deploy congruent mindreading instead of non-congruent mindreading, the core semantic properties of confirmatives should be easier to learn than *think* because it is doxastic and yet non-veridical. But an obstacle to investigating as much is the apparent infrequency of these verbs. For example, the Gleason corpus in CHILDES (Masur and Gleason, 1980; MacWhinney, 2000), which is the corpus Dudley et al. (2017) investigate, contains only a single instance of *prove*, zero instances of *confirm*, and no instances of *show* with a *that*-complement.

Another way forward is through assumptions children make about veridical verbs. If children assume veridical verbs are associated with a property F, then whether children regard *think* as veridical can be investigated by considering whether children regard *think* as F too. One candidate is a verb's ability to embed both declarative and interrogative complements, a property that is known as RESPONSIVITY (Lahiri, 2002). Confirmative verbs appear to be consistently responsive. (14) and (15) illustrate.

(14) The professor proved what $\sqrt{400}$ is.

(15) Margaret confirmed who took a bath.

In contrast, doxastic verbs like *think* are widely regarded as non-responsive. However, this generalization and nearby variants are false. As White (2021) shows with large-scale experimental data and instances found in corpora, *think* can be responsive. Here are some of his examples.

(16) When Jan Brown completed her safety briefing for the passengers, she tried to think whether she had covered everything.

(17) I'm trying to think whether I'd have been a star today or not.

Nevertheless, children plausibly operate with some version of this generalization. One reason to think this is that the high-frequency verbs, the verbs children most encounter, are verbs where the generalization does hold (White and Rawlins, 2018). In sorting verbs into different semantic classes, we have reason to expect that children are regarding veridical verbs as responsive verbs. As such, the veridicality hypothesis can be evaluated by exploring to extent to which three-year-old children regard *think* as responsive. Finding that children regard *think* as more responsive than adults do would provide evidence that children are treating *think* as being akin to a confirmative verb.

Whether the veridicality hypothesis is correct therefore remains an open question. But this essay has attempted to show that it is a live option that merits consideration, and that it recommends fruitful topics to investigate in connection to the acquisition of *think*.[8]

Conflict of interest

There are no conflicts of interest to declare.

References

Leonard Abbeduto and Sheldon Rosenberg. Children's knowledge of the presuppositions of *know* and other cognitive verbs. *Journal of Child Language*, 12 (3):621–641, 1985.

Marie Aguirre, Mélanie Brun, Anna Reboul, and Oliver Mascaro. How do we interpret questions? Simplified representations of knowledge guide humans' interpretation of information requests. *Cognition*, 218:104954, 2022.

Pranav Anand and Valentine Hacquard. Factivity, belief and discourse. In L. Crnic and U. Sauerland, editors, *The Art and Craft of Semantics: A Festschrift for Irene Heim*, pages 69–90. MITWPL, 2014.

[8] I am indebted to two anonymous referees, the associate editor, Patrick G. Grosz, and the editors-in-chief, Michael Glanzberg and Yael Sharvit, for feedback on earlier drafts of this paper.

Scott AnderBois. On the exceptional status of reportative evidentials. *Proceedings of SALT*, 24:234–254, 2004. doi: https://doi.org/10.3765/salt.v24i0.2424.

Mathew Benton and Peter van Elswyk. Hedged assertion. In Sandy Goldberg, editor, *Oxford Handbook of Assertion*, pages 245–263. Oxford University Press, 2020.

Matthew Benton. Two more for knowledge account of assertion. *Analysis*, 71:684–687, 2011.

Matthew Benton. Gricean quality. *Noûs*, 50(4):689–703, 2016.

Matthew Benton. Knowledge is the norm of assertion. In B. Roeber, J. Turri, M. Steup, and E. Sosa, editors, *Contemporary Debates in Epistemology*. Blackwell, 3rd edition, forthcoming.

Lois Bloom, Matthew Rispoli, Barbara Gartner, and Jeremie Hafitz. Acquisition of complementation. *Journal of Child Language*, 16(1):101–120, 1989.

Paul Bloom. Mindreading, communication and the learning of names for things. *Mind and Language*, 17(1-2):37–54, 2002.

Sarah Brown-Schmidt and Scott Fraundorf. Interpretation of informational questions modulated by joint knowledge and intonational contours. *Journal of Memory and Language*, 84:49–74, 2015.

David Buttelmann, Malinda Carpenter, and Michael Tomasello. Eighteen-month-old infants show false belief understanding in an active helping paradigm. *Cognition*, 112(2):337–342, 2009.

Josep Call and Michael Tomasello. A nonverbal false belief task: The performance of children and great apes. *Child Development*, 70(2):381–395, 1999.

Susan Carey. The child as word learner. In M. Halle, J. Bresnan, and G. Miller, editors, *Linguistic theory and psychological reality*, pages 264–293. MIT Press, 1978.

Gennaro Chierchia and S. McConnell-Ginet. *Meaning and Grammar: An Introduction ot Semantics*. MIT Press, 2nd edition, 2000.

Kathleen Corriveau, Paul Harris, Elizabeth Meins, Charles Fernyhough, Bronia Arnott, Lorna Elliott, Beth Liddle, Alexandra Hearn, Lucia Vittorini, and Marc de Rosnay. Young children's trust in their mother's claims: Longitudinal links with attachment security in infancy. *Child Development*, 80(3):750–761, 2009.

Jill de Villiers and Jennie Pyers. Complements to cognition: A longitudinal study of the relationship between complex syntax and false-belief-understanding. *Cognitive Development*, 17(1):1037–1060, 2002.

Holger Diessel and Michael Tomasello. The acquisition of finite complement clauses in english: A corpus-based analysis. *Cognitive Linguistics*, 12(2):97–141, 2001.

Ceymi Doenyas, H. Melis Yavuz, and Bilge Selcuk. Not just a sum of its parts: How tasks of the theory of mind scale relate to executive function across time. *Journal Of Experimental Child Psychology*, 166:485–501, 2018.

Rachel Dudley, Naho Orita, Valentine Hacquard, and Jeffrey Lidz. Three-year-olds' understanding of *know* and *think*. In Florian Schwarz, editor, *Experimental perspectives on presuppositions*, pages 241–262. Springer, 2015.

Rachel Dudley, Meredith Rowe, Valentine Hacquard, and Jeffrey Lidz. Discovering the factivity of *know*. *Semantics and Linguistics Theory*, 27:600–619, 2017.

Paul Egré. Question-embedding and factivity. *Grazer Philosophische Studien*, 77(1): 85–2008, 2008.

Martina Faller. *Semantics and Pragmatics of Evidentials in Cuzco Quechua*. PhD thesis, Stanford University, 2002.

Rachel Joffe Falmagne, Joanna Gonsalves, and Sarah Bennett-Lau. Children's linguistic intuitions about factive presuppositions. *Cognitive Development*, 9:1–22, 1994.

Jane Gillette, Henry Gleitman, Lila Gleitman, and Anne Lederer. Human simulations of vocabulary learning. *Cognition*, 72(2):135–176, 1999.

Lila Gleitman. The structural sources of verb meanings. *Language Acquisition*, 1(1): 3–55, 1990.

Valentine Hacquard and Jeffrey Lidz. Children's attitude problems; bootstrapping verb meaning from syntax and pragmatics. *Mind and Language*, 34(1):73–96, 2019.

Valentine Hacquard and Jeffrey Lidz. On the acquisition of attitude verbs. *Annaul Review of Linguistics*, 8:193–212, 2022.

J. Kiley Hamlin, Tomer Ullman, Josh Tenenbaum, Noah Goodman, and Chris Baker. The mentalistic basis of core social cognition: experiments in preverbal infants and a computational model. *Developmental Science*, 16(2):209–226, 2013.

Kaitlyn Harrigan, Valentine Hacquard, and Jeffrey Lidz. *Hope* for syntactic bootstrapping. *Language*, 95(4):642–82, 2019.

Kate Harrigan, Valentine Hacquard, and Jeffrey Lidz. Three-years-olds' understanding of desire reports is robust to conflict. *Frontiers in Psychology*, 9:119, 2018.

Paul Harris. Trust. *Developmental Science*, 10(1):135–138, 2007.

Tone Hermansen, Samuel Ronfard, Paul Harris, Francisco Pons, and Imac Zambrana. Young children update their trust in an informant's claim when experience tells them otherwise. *Journal Of Experimental Child Psychology*, 205:105063, 2021.

Joan B. Hooper. On assertive predicates. In John P. Kimball, editor, *Syntax and Semantics 4*, pages 91–123. Academic Press, 1975.

Nick Huang, Aaron White, Chia-Hsuan Liao, Valentine Hacquard, and Jeffrey Lidz. Syntactic bootstrapping attitude verbs despite impoverished morphosyntax. *Language Acquisition*, 29(1):27–53, 2022.

Yi Ting Huang and Jesse Snedeker. Semantic meaning and pragmatic interpretation in 5-year-olds: Evidence from real-time spoken language comprehension. *Developmemental Psychology*, 45(6):1723–1739, 2009.

Hande Ilgaz, Jedediah Wilfred Papas Allen, and Feride Nur Haskaraca. Is cultural variation the norm? A closer look at sequencing of the theory of mind scale. *Cognitive Development*, 63:101216, 2022.

Carl Nils Johnson and Michael Maratsos. Early comprehension of mental verbs: *think* and *know*. *Child Development*, 48(4):1743–1747, 1977.

Marina Kammermeier and Markus Paulus. Do action-based tasks evidence false-belief understanding in young children? *Cognitive Development*, 46:31–39, 2018.

Lauri Kartunnen. Implicative verbs. *Language*, 47(2):340–358, 1971.

Napoleon Katsos, Chris Cummins, Maria-José Ezeizabarrena, Anna Gavarró, Jelena Kuvac Kraljevic, Gordana Hrzico, Kleanthese Grohmann, Athina Skordi,

Kristine Jensen de López, Lone Sundahl, Angeliek van Hout, Bart Hollebrandse, Jessica Overweg, Myrthe Faber, Margreet van Koert, Nafsika Smith, Maigi Vija, Sirli Zupping, Sari Kunnari, Tiffany Morisseau, Manana Rusieshvili, Kazuko Yatsushiro, Anja Fengler, Spyridoula Varlokosta, Katerina Konstantzou, Shira Farby, Maria Teresa Guasti, Mirta Vernice, Reiko Okabe, Miwa Isobe, Peter Crosthwaite, Yoonjee Hong, Inrida Balciuniene, Yanti Marina Ahmad Nizar, Helen Grech, Daniela Gatt, Win Nee Cheong, Arve Asbjornsen, Janne von Koss Torkildsen, Ewa Haman, Aneta Miekisz, Natalia Gagarina, Julia Puzanova, Darinka Andelkovic, Maja Savic, Smiljana Josic, Daniela Slancova, Svetlana Kapalkova, Tania Barberan, Duygo Ozge, Saima Hassan, Cecilia Yuet Hung Chan, Tomoya Okubo, Heather van der Lely, Uli Sauerland, and Ira Noveck. Cross-linguistic patterns in the acquisition of quantifiers. *Proceedings of the National Academy of Sciences of the United States of America*, 113(33):9244–9249, 2016.

Melanie Khu, Craig Chambers, and Susan Graham. Preschoolers flexibly shift between speakers' perspectives during real-time language comprehension. *Child Development*, 91(3):619–634, 2020.

Peter Kinderman, Robin Dunbar, and Richard Bentall. Theory-of-mind deficits and causal attributions. *British Journal of Psychology*, 89(2):191–204, 1998.

Carla Krachun, Malinda Carpenter, Josep Call, and Michael Tomasello. A competitive nonverbal false belief task for children and apes. *Developmental Science*, 12(4):521–535, 2009.

Utpal Lahiri. *Questions and answers in embedded contexts*. Oxford University Press, 2002.

Barbara Landau and Lila Gleitman. *Language and experience: Evidence from the blind child*. Harvard University Press, 2009.

Jacques Launay, Eiluned Pearce, Rafael Wlodarski, Max van Duijn, and James Carney. Higher-order mentalizing and executive functioning. *Personality and individual differences*, 86:6–14, 2015.

Alan Leslie, Ori Friedman, and Tim German. Core mechanisms in "theory of mind". *Trends in Cognitive Science*, 8(12):528–533, 2004.

Shevaun Lewis, Valentine Hacquard, and Jeffrey Lidz. "think" pragmatically: Children's interpretation of belief reports. *Language Learning and Development*, 13 (4):395–417, 2017.

Bethany Liddle and Daniel Nettle. Higher-order theory of mind and social competence in school-age children. *Journal of Cultural and Evolutionary Psychology*, 43(3):1489–5254, 2006.

Yuyan Luo and Renée Baillargeon. Do 12.5-month-old infants consider what objects others can see when interpreting their actions? *Cognition*, 105(3):489–512, 2007.

Yuyan Luo and Susan Johnson. Recognizing the role of perception in action at 6 months. *Developmental Science*, 12(1):142–149, 2009.

Brian MacWhinney. *The CHILDES project: Tools for analyzing talk*. Lawrence Erlbaum, 3rd edition, 2000.

Elise Masur and Jean Gleason. Parent-child interaction and the acquisition of lexical information during play. *Developmental Psychology*, 16(5):404, 1980.

Chris Moore and Jane Davidge. The development of mental terms: pragmatics or semantics? *Journal of Child Language*, 16(3):633–641, 1989.

Chris Moore, Dana Bryant, and David Furrow. Mental terms and the development of certainty. *Child Development*, 60(1):167–171, 1989.

Sarah Murray. *The Semantics of Evidentials*. Oxford: Oxford University Press, 2017.

Aparna Nadig and Julie Sedivy. Evidence of perspective-taking constraints on children's on-line reference resolution. *Psychological Science*, 13(4):329–336, 2002.

Jennifer Nagel. Knowledge as a mental state. *Oxford Studies in Epistemology*, 4:275–310, 2013.

Jennifer Nagel. Factive and non-factive mental state attribution. *Mind and Language*, 32:525–544, 2017.

Letitia Naigles. Manipulating the input: Studies in mental verb acquisition. In Barbara Landau, John Sabini, John Jonides, and Elissa Newport, editors, *Perception, cognition, and language: Essays in honor of Henry and Lila Gleitman*, pages 245–274. MIT Press, 2000.

Ira Noveck. When children are more logical than adults: experimental investigations of scalar implicatures. *Cognition*, 78(2):165–188, 2001.

Nathan Oesch and Robin Dunbar. The emergence of recursion in human language: mentalizing predicts recursive syntax task performance. *Journal of Neurolinguistics*, 43:95–106, 2017.

Kristine Onishi and Renée Baillargeon. Do 15-month-old infants understand false beliefs? *Science*, 308(5719):255–258, 2005.

Anna Papafragou. Epistemic modality and truth conditions. *Lingua*, 116:1688–1702, 2006.

Anna Papafragou and Julien Musolino. Scalar implicatures: Experiments at the semantics-pragmatics interface. *Cognition*, 86(3):253–282, 2003.

Josef Perner, Manuel Sprung, Petra Zauner, and Hubert Haider. *Want that* is understood well before *say that*, *think that*, and false belief. *Child Development*, 74 (1):179–188, 2003.

Joseph Perner. On representing that: The asymmetry between belief and desire in 38 children's theory of mind. In D. Frye and C. Moore, editors, *Children's theories of mind: Mental states and social understanding*, pages 139–156. Lawrene Erlbaum, 1991.

Jonathan Phillips and Aaron Norby. Factive theory of mind. *Mind and Language*, 36: 3–26, 2021.

Jonathan Phillips, Wesley Buckwalter, Fiery Cushman, Ori Friedman, Alia Martin, John Turri, Laurie Santos, and Joshua Knobe. Knowledge before belief. *Behavioral and Brain Sciences*, 44:1–75, 2021.

Diane Poulin-Dubois, Hannes Rakoczy, Kimberly Burnside, Cristina Crivello, Sebastian Dörrenberg, Katheryn Edwards, Horst Krist, Louisa Kukle, Ulf Liszkowski, Jason Low, Josef Perner, Lindsey Powell, Beate Priewasser, Eva Rafetseder, and Ted Ruffman. Do infants understand false beliefs? We don't know yet. *Cognitive Development*, 48:302–215, 2018.

Lindsey Powell, Kathryn Hobbs, Alexandros Bardis, Susan Carey, and Rebecca Saxe. Replications of implicit theory of mind tasks with varying representational demands. *Cognitive Development*, 46:40–50, 2018.

William Van Orman Quine. *Word and Object*. MIT Press, 1960.

Samuel Ronfard, Evan Chen, and Paul Harris. The emergence of the empirical stance: Children's testing of counterintuitive claims. *Developmemental Psychology*, 54(3): 482–493, 2017.

Jason Scofield and Douglas Behrend. Learning words from reliable and unreliable speakers. *Cognitive Development*, 23:278–290, 2008.

Ameneh Shahaeian, Candida Peterson, Virginia Slaughter, and Henry Wellman. Culture and the sequence of steps in theory of mind development. *Developmemental Psychology*, 47(5):1239–1247, 2011.

Mandy Simons. Observations on embedding verbs, evidentiality, and presupposition. *Lingua*, 117(6):1034–1056, 2007.

Beate Sodian, Claudia Thoermer, and Ulrike Metz. Now I see it but you don't: 14-mont-olds can represent another person's visual perspective. *Devleopmental Science*, 10(2):F1–F14, 2007.

Victoria Southgate, Atsushi Senju, and Gergeley Csibra. Action anticipation through attribution of false beliefs by 2-year-olds. *Psychological Science*, 18(7):587–592, 2007.

Benjamine Spector and Paul Egré. A uniform semantics for embedded interrogatives: an answer, not necessarily the answer. *Synthese*, 192:1729–1784, 2015.

Asbjørn Steglich-Petersen and John Michael. Why desire reasoning is developmentally prior to belief reasoning. *Mind and Language*, 30(5):526–549, 2015.

Luca Surian, Stefania Caldi, and Dan Sperber. Attribution of beliefs by 13-month-old infants. *Psychological Science*, 18(7):580–586, 2007.

Twila Tardif and Henry Wellman. Acquisition of mental state language in mandarin- and cantonese-speaking children. *Developmemental Psychology*, 36(1):25–43, 2000.

Birgit Träuble, Vesna Marinović, and Sabina Pauen. Early theory of mind competencies: Do infants understand others' beliefs? *Infancy*, 5(4):434–444, 2010.

John Turri. Prompting challenges. *Analysis*, 70:456–462, 2010. doi: https://doi.org/10.1093/analys/anq027.

John Turri. The point of assertion is to transmit knowledge. *Analysis*, 76(2):130–136, 2016.

John Turri. Experimental work on the norms of assertion. *Philosophy Compass*, 12(7): e12425, 2017. doi: https://doi.org/10.1111/phc3.12425.

John Turri. Revisiting norms of assertion. *Cognition*, 177:8–11, 2018. doi: http://doi.org/10.1016/j.cognition.2018.03.023.

J.O. Urmson. Parenthetical verbs. *Mind*, 61(244):480–496, 1952.

Peter van Elswyk. Hedged testimony. *Noûs*, 57(2):341–369, 2023.

Peter van Elswyk. Hedging in discourse. *Synthese*, 204(98):1–31, 2024.

Peter van Elswyk and Matthew Benton. Assertion remains strong. *Philosophical Studies*, 180:27–50, 2023.

Henry Wellman and David Liu. Scaling of theory-of-mind tasks. *Child Development*, 75(2):523–541, 2004.

Henry Wellman and Jacqueline Woolley. From simple desires to ordinary beliefs: The early development of everyday psychology. *Cognition*, 35(3):245–275, 1990.

Henry Wellman, David Cross, and Julanne Watson. Meta-analysis of theory-of-mind development: the truth about false belief. *Child Development*, 72(3):655–684, 2001.

Evan Westra and Peter Carruthers. Pragmatic development expains the theory-of-mind scale. *Cognition*, 158(C):165–176, 2017.

Evan Westra and Jennifer Nagel. Mindreading in conversation. *Cognition*, 210 (104618):1–15, 2021.

Aaron White. On believing and hoping whether. *Semantics and Pragmatics*, 14(6): 1–18, 2021.

Andrew White and Kyle Rawlins. The role of veridicality and factivity in clause selection. In *48th Annual Meeting of the North East Linguistic Society*, pages 1–14, 2018.

Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, 2000.

H. Wimmer and J. Perner. Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 130(1):128–133, 1983.

Jessica Yott and Diane Poulin-Dubois. Are infants' theory-of-mind abilities well integrated? implicit understanding of intentions, desires, and beliefs. *Journal of Cognition and Development*, 17(6):683–698, 2016.

Karolina Zuchewicz. *On the veridicality of perfective clause-embedding verbs in Polish: A unified aspect-based analysis of incremental theme verbs with nominal and propositional complements*. PhD thesis, Humboldt-Universität zu Berlin, 2020.