# Knowledge, the concept KNOW, and the word *know*: considerations from polysemy and pragmatics

Rachel Dudley
*Central European University*

Christopher Vogel
*Shandong University*

**Abstract** A recent focus on philosophical methodology has reinvigorated ordinary language philosophy with the contention that philosophical inquiry is better served by attending to the ordinary use of language. Taking cues from findings in the social sciences that deploy methods utilizing language, various ordinary language philosophers embrace a guiding mandate: that ordinary language usage is more reflective of our linguistic and conceptual competencies than standard philosophical methods. We analyze two hypotheses that are implicit in the research from which ordinary language approaches take their cues. This pair of optimistic assumptions (a) bind word meanings to properties of their corresponding concepts and (b) regard language as a direct reflection of our underlying cognitive processes and competencies. Polysemy and pragmatics complicate each assumption. Because the ordinary language philosopher's methodological mandate compels us to consider how individuals process the utterances they encounter in deciphering the communicative intentions of speakers, failing to attend to the import of polysemy and pragmatics in philosophical and empirical methods has the potential to frustrate the aims of their insightful mandate. The significance of those two complications is worked out with the case study of knowledge.

Penultimate draft of:
Dudley, Rachel & Christopher A. Vogel. 2024. Knowledge, the concept KNOW and the word *know*: considerations from polysemy and pragmatics. *Synthese* 203(17)
DOI: https://doi.org/10.1007/s11229-023-04439-1
*Please cite published version.*

For helpful comments on versions of this paper we thank three anonymous reviewers at this journal, as well as the participants at the conference 'Philosophy's Experimental Turn and the Challenge from Ordinary Language' (Berlin, June 2022). Special thanks is also owed to Eugen Fischer and Nat Hansen for their insightful input at various stages of this project. This project is funded with support from the European Union's Horizon 2020 research and innovation programme via a Widening Fellowship #101038068 (ISLI).

## 1 Introduction

A recent focus on philosophical methodology has reinvigorated ordinary language philosophy (Hansen 2020, Cappelen & McKeever 2022). The contemporary incarnation of this methodological camp contends that philosophical inquiry is better served by attending to the ordinary use of language. This has manifested in a number of ways, many of which take cues from methods (and findings) in the social sciences (see for example Hansen & Chemla 2015, Semler & Henne 2019, Ventura 2019, Chartrand 2022). Because the guiding principle of ordinary language philosophy insists that ordinary language usage is more reflective of our linguistic and conceptual competence, these interdisciplinary sources invoke methods that utilize language. This utilization of linguistic data to inform the object of study occurs elsewhere throughout cognitive science, often because language offers both a quick shorthand for communicating stimuli to linguistic minds and for requesting responses from them.

Such cross-disciplinary interaction should be welcomed, given the overlap in what matters to philosophers, psychologists, and linguists. However, as we argue, more attention needs to be paid to the role of linguistic and semantic hypotheses in utilizing these methods, and in synthesizing cross-disciplinary findings. We argue that the following pair of **optimistic assumptions** motivate findings in much of philosophy and cognitive science: (1) that natural language expressions have meanings that determine truth conditions, and (2) that a speaker's utterances are straightforwardly reflective of their underlying linguistic and/or extra-linguistic cognition. These related assumptions are in tension with facts about natural language meanings, on the one hand, and the variety of pragmatic factors involved in interpreting language on the other. Our purpose is not to show that these assumptions are false—no single paper could hope to achieve that aim. Rather, we point to why these assumptions deserve to be regarded with more skepticism, why they obscure the results of inquiry, and why a commitment to ordinary language philosophy is discordant with them.

In what follows we make the case for heeding this methodological prescription by focusing on a particular case of significant interdisciplinary overlap between philosophers, psychologists, developmentalists, linguists, and cognitive scientists generally: the case of knowledge. The overlap—between these various projects that study (some combination of) linguistic expressions like *know*, concepts like KNOW, and states of knowledge[1]—offers fertile ground for cross-disciplinary interaction.

1 For the sake of clarity we adopt the following orthographic convention in what follows: when discussing linguistic expressions we mark them with italics (e.g., *know*); when discussing concepts we use the small-caps (e.g., KNOW); when talking about worldly objects we leave the relevant expressions unmodified.

But precisely because of this fertility, and the central role that language plays in the methods utilized throughout a shared course of inquiry, attending to the centrality of the optimistic assumptions is essential for making use of the findings and methods to advance within each discipline.

Before turning to the case of knowledge as exemplar, we begin with some preliminary assumptions about language and concepts (§2). We continue by outlining the central problem with the method of cases that the current revival of ordinary language philosophy emerged to resolve, the **productive limit problem** (§3.1). We then briefly survey the various solutions offered by ordinary language philosophy, to illustrate that none of them address a more **fundamental problem** with the method of cases (§3.2). Rather this problem is often resolved by the first **optimistic assumption** that meanings determine intensions/concepts/satisfaction-conditions. We argue that this assumption is in tension with a shared commitment of (most) natural language philosophers, a tension borne out by the **polysemy** of natural language expressions (§3.3).

Natural language polysemy, that words bear a one-to-many relation with concepts/ intensions/satisfaction-conditions, mirrors a related phenomena in language use, in that a speaker can use a word with a variety of communicative purposes, with a one-to-many relation obtaining between words and intentions/purposes. We argue that this **pragmatic** aspect of language use is often ignored in empirical work regarding both language and extra-linguistic cognition, under the second **optimistic assumption** that a participant's set of behaviors directly reflects just the desired target of study.

To show this, we review studies from developmental psychology (§4) that regard the properties of the mind-reading system responsible for knowledge-attribution (§4.2). We then show that central findings regarding the development of children's mind-reading (or Theory of Mind) capacities admit to pragmatic confounds (§4.3). To illuminate these pragmatic challenges further, we then turn to the study of the linguistic expression *know* (§5), detailing features of the verb (§5.1). We then review studies from the language acquisition literature to demonstrate how children's understanding of propositional attitude reports is impacted by pragmatic principles (§5.2.1), and furthermore suggest that this is not surprising, given that such pragmatic principles are operative in how children acquire the semantics of these attitude reports in the first place (§5.2.2). We conclude with some final thoughts (§6).

## 2 Preliminaries on language, concepts and interpretation

Our central aim in what follows is to motivate a move toward greater attention to the role that the **optimistic assumptions** about language play in inquiry across disciplines. For that reason, we flag here a set of assumptions that we will make

in the following sections. The first is about the operative conception of language utilized throughout. The second is a related clarification about concepts, and the third about the relationship between language and what it is used to achieve.

We take as a starting point that children around the age of 4 years exhibit linguistic competence of a kind that human adults enjoy. Children develop, or in some way internalize, procedures to generate linguistic expressions that adults in their community understand. This suggests, prima facie, two options for what the study of language is the study of: the **expressions** that speakers generate, or the **procedures** responsible for generating those expressions.

The first approach takes languages to be mind-external objects, for example sets of expression-meaning pairs (Lewis 1975). As such, children acquire a language by developing representations of this abstract object. Following custom, we can call such objects **E-languages**. The second notion of language concerns the procedures that generate expressions—the mind-internal system(s) young children develop (and adults have) which are utilized in forming and understanding expressions. Call these **I-languages**. The following claim then seems obviously true:

(1) I-languages generate E-languages.

But such a conclusion would be hasty. After all, (1) offers an empirical hypothesis, that (at the very least) assumes there are both I-languages and E-languages. But neither of these kinds of thing exist merely as a consequence of offering a definition—much less by way of the loose demonstration above. So, do either I-languages or E-languages exist?

Let's start with I-languages. Children who gain linguistic competence must somehow develop procedures for constructing the expressions that adults understand. To call such procedures I-languages assumes minimally that children in the same linguistic community develop similar procedures. But such similarity does not merely follow from recognizing that children develop procedures to generate comprehensible linguistic expressions simpliciter. Individual children may well develop vastly different procedures for generating the same collection of comprehensible expressions, in much the way that different geometry students might follow different steps to aptly bisect the angles of the same square (e.g., one utilizing a compass, another a protractor, and a third only a straight edge). If different children develop vastly different procedures for generating expressions, there are no (cohesive) I-languages.

Similarly, given cohesive I-languages, it does not immediately follow that there are E-languages. That is, the expressions produced by I-language procedures may not form a cohesive class that can be delineated independently of the procedures that generated them. Take for example a simple proposal about E-languages that takes them to be sets of sentences; the object so-called "English" is just an infinitely

long list of expressions. That is, there is a set of things, like the following, that some children triangulate on in learning a language:

(2) a. The photographer stole the car with a screwdriver.
b. The photographer stole the car with a nice paint job.
c. The photographer started the car with a screwdriver.
d. The photographer stole the car that had a screwdriver in the trunk.
e. The photographer stole the car.
. . .

But we might worry about our ability to individuate sentences that populate this set, given that (2a) illustrates the constrained homophony of natural language expressions. This one expression has at least two meanings. More precisely, an English speaker can use it to talk about how someone used a screwdriver to start and then take a car (in the way (2c) expresses), or to describe someone that took a car that had a screwdriver in the trunk (in the way (2d) expresses). Since this difference between the uses of (2a) depends on the underlying grammatical structure of the expression, one might worry whether the expression in (2a) indicates one sentence or two. If (2a) really indicates two different sentences, how does one go about differentiating them, without appealing to the generative procedures (i.e., grammars) that produce them? Further difficulties arise in positing and individuating E-languages, ranging from considerations of pronunciation and phonology (Bromberger 2011), to diachronic semantic drift (Bolinger 2020) to theoretical utility (Chomsky 2000).[2]

Following Chomsky (1986), we take the study of language as primarily concerned with the internalized, and intensionally individuated procedures responsible for generating expressions. One central reason for taking languages to be aspects of human minds has to do with language acquisition, and poverty of the stimulus arguments (Chomsky 1965, 1980, Crain & Pietroski 2001, Berwick et al. 2011). But, more importantly, if one takes languages to be E-languages it is difficult to see how children could ever come to learn languages, let alone with a rapidity that outpaces their input. How does a child come to (partially, cf. Dummett 1978, 1986, 1993) learn the members of a (infinite) set that share no obvious sui generis properties?

A related preliminary point pertains to our discussion of concepts. The term "concept" gets defined in a multitude of ways in philosophy, psychology, linguistics, sociology, and the social sciences more broadly. So it will be fruitful to clarify how we intend it here. Because we take the discussion of linguistic expressions to pertain to I-languages, which are cognitive systems, we likewise take concepts to be mental resources. Linguistic meaning, after all, is related to thought. Put loosely, words and their meanings are somehow related to concepts. As such, if linguistic expressions

2 For a more exhaustive discussion see Ludlow (2003). For a defense of E-languages see Devitt (2006).

are the product of procedures for generating complex linguistic representations from simpler ones (e.g., generating sentences from words), their meanings must bear some relation to the concepts utilized in extra-linguistic cognition. Thus, we take concepts to be mental resources (and not say, mind-independent objects; see Zalta 2001 for discussion), without specifying the precise nature of the conceptual system.

Lastly, taking I-languages as the core competence that individuals develop in language acquisition still leaves unattested the variability we see in communication. Speakers utilize this competence in generating and externalizing expressions in order to communicate, but do so with many purposes or intentions. Just as we might worry about our ability to individuate the sentences in a set like (2), we should also worry about how to individuate the different intentions that a speaker could have in uttering any of the members of this set. For example, a sentence like (2e) could be uttered or written by a journalist in order to inform the public about a series of events that occurred. But, when uttered by a detective, it can ultimately function as a directive to a subordinate to arrest a particular criminal (Austin 1975, Searle 1975). As researchers, this bears on the task of uncovering aspects of linguistic and conceptual competence, since the same pragmatic abilities utilized by listeners to discern what a speaker intends to communicate by a given utterance are also utilized by speakers in producing the expressions that convey their communicative intentions. Since research participants are making use of this pragmatic competence to communicate with investigators in responding to linguistic prompts, as we'll argue, accounting for this competence is essential since utterances offer only indirect evidence about both the I-language procedures that generate them and the mental representations they recruit in extra-linguistic cognition.

One final caveat regarding what follows. Our contention is that more attention should be paid to the role of the linguistic and semantic hypotheses at work in cross-disciplinary efforts. In the exemplar case of the study of knowledge-related phenomena we characterize this in terms of the relations between expressions like *know*, concepts like KNOW and mental states of knowing.[3] Doing so abstracts away from many details that practitioners in various disciplines would find troubling. Take words for example. Many who take the notion of I-languages seriously claim that there is no coherent way to delineate what a word is. Rather, we should appeal to phonological and syntactic features, and interfacing relations with the conceptual system (see Clark 2020 for a helpful discussion). Further elaborations of this sort are within the spirit of what follows, which is described at a level of abstraction

3 A recent Knowledge First approach in epistemology argues that knowledge is a purely mental state, in contrast to the orthodox analysis of knowledge as belief-states that meet some further, often mind-external condition. Our talk of "knowing as a mental state" does not to prejudge such debates, but simply indicates the object that those concerned with cognitive ontology are investigating, irrespective of how that thing, whatever it turns out to be, relates to words like *know* or concepts like KNOW.

removed from such disagreements, as to make the argument digestible to a broad interdisciplinary audience.

## 3 Studying knowledge

Our central aim of underscoring the optimistic assumptions implicit in methods that utilize language begins with the study of knowledge. As often investigated by epistemologists, the study of knowledge makes use of the **method of cases**. In §3.1 we briefly describe that method and the role of language therein, in the service of articulating the **productive limit problem** that much ordinary language philosophy aims to resolve. We argue that there is a more **fundamental problem** with the method of cases, and show in §3.2 that the solutions offered by ordinary language philosophy do not address this more central concern. In §3.3 we argue that straightforward ways of resolving the fundamental problem are undermined by the pervasive **polysemy** of natural language expressions, and further impose a burden on natural language philosophers that is in tension with their appeals to the richness of ordinary language use in mending the method of cases.

### 3.1 The method of cases and the productive limit problem

The most direct study of the nature of knowledge itself is the **analysis of knowledge**, which strives to offer the conditions that a mental state needs to satisfy in order to be a knowledge state.[4] The **orthodox** analysis of knowledge treats knowledge as a species of belief. Knowing is a matter of having a belief that is both true and justified. Gettier (1963) famously offers counter-examples to the orthodox account of knowledge as justified true belief, describing a Smith who finds himself in a position of epistemic luck. Utilizing the **method of cases** he offers a short **Vignette** describing a situation in which a job-seeking Smith, given solid evidence and a bit of logic, comes to believe that "(e) The man who will get the job has ten coins in his pocket" (Gettier 1963: pg. 122). Assuming that a competitor will get the job, Smith is surprised to learn that not only did he get the job, but that he had ten coins in his pocket, making his belief both true and (seemingly) justified.

Readers are then offered a **Prompt**, invited to accept or reject the claim that "Smith does not know that (e) is true" (Gettier 1963: pg. 122). The lesson from Gettier's case, accepted almost universally among philosophers and epistemologists

4 The primacy of intuitions as evidence in investigating knowledge is not universally taken as central, or even necessary (Kornblith 2002, Maffie 1990). Kornblith (2002) for example takes knowledge to be natural kind (*ibid* §2.7), which our intuitive judgments might reliably track, but no more requires the method of cases than other naturalistic investigation (*ibid* §1.3). See Goldman (2005) for discussion. Our thanks to an anonymous reviewer for drawing our attention to this point.

(Starmans & Friedman 2020), is that, to count as knowledge, a justified true belief must meet some further condition. The ensuing decades of work on the analysis of knowledge can be fairly characterized as a series of attempted proposals, cases, and counterexamples regarding what could save the orthodoxy (Shope 1983).

Common to this history is the use of a standard philosophical methodology, the method of cases. The familiar method solicits judgments from readers regarding a creative **Vignette** in order to probe the target of inquiry with a **Prompt**. Gettier's paper is a classic case of the standard methodology in practice. When the method is (purportedly) successful, audience judgments/intuitions are univocal, and thereby yield verdicts about the subject of inquiry.[5]

In the case of the analysis of knowledge, the method of cases seems to have reached its productive limit. As novel cases get generated to refute proposals about the proper analysis of knowledge, new views get proposed, only to suffer from further (more elaborate) counter-examples. Such efforts seem doomed to fail, since Gettier cases can be generated for seemingly any proposal, given a few highly plausible assumptions (Zagzebski 1994). Some philosophers have taken such failure to indicate that something is amiss with the task of analyzing knowledge (Williamson 2000, Nagel 2013). But the problem seems more fundamental, reflecting a problem, not with the orthodox project of analyzing knowledge, but with the method of cases (Baz 2017, Vogel 2018). Call this the **productive limit problem**.

A contemporary thread in response to this problem concerns the role of language in the method of cases. These **ordinary language philosophers** take the use of language in the method of cases as being, in one way or another, not genuinely reflective of an audience's competence, thereby distorting the results of inquiry. One can glean from the literature four general ordinary language philosophy strategies for resolving the productive limit problem. We focus on epistemology and the analysis of knowledge here, but these strategies can (and have) been applied more broadly to other areas (see, for example, Li 2022 for a discussion of similar efforts regarding proper names that emerged in response to Machery et al. 2004).

What we argue in this section is that, while this focus on language in responding to the productive limit problem is apt, the four kinds of solutions offered miss a more **fundamental problem** with the method of cases. The fundamental problem is that, in the absence of ancillary (optimistic) assumptions, the method of cases produces equivocal results. As we will argue in §3.3, the most straightforward ways of addressing this problem are unavailable to the ordinary language philosopher.

5 Gettier's paper deploys the method of cases, but does so in a way that superficially deviates from our presentation here. After offering a vignette about Jones and Smith, Gettier simply *states* that Smith does not have knowledge: "Smith does not know that (e) is true" (Gettier 1963: pg. 122). However, the implication is clear, that Gettier's reader would agree with him.

A contrast with the use of the method of cases will be helpful in making the fundamental problem salient. Linguists often use **acceptability judgment tasks**, where participants are provided linguistic inputs, and then asked for their intuitive judgments about those stimuli. The methodology is used to explore the different kinds of strings that grammars generate. As practiced from the armchair, this method is utilized in written work whereby the author presents the reader with expressions, and simply announces a verdict about their (un)acceptability. As with the method of cases, such proclamations implicitly invite the reader to assess the expressions themselves, with the assumption that they will agree. Commonly, a **minimal pair** is given to make the point, with the unacceptable one of the pair marked with an asterisk. For example, in the following pair, (3) seems perfectly acceptable, while (4) is not:

(3) The collar of the dog that I saw was leather.

(4) * The dog that I saw's collar was leather.

This serves as a data point pertaining to the grammar of possessives with complex noun phrases, since the sentence in (5) seems perfectly acceptable.

(5) Toto's collar was leather.

The readers' (and author's) judgments count as data to be explained by some theory of the grammar of possessives, granting the assumption that acceptability judgments are driven by cognitive systems that encode grammar (see Marinis 2016 for discussion). Thus, the sentence in (4) is unacceptable to English speakers because it is ungrammatical.

Much like the method of cases, a given linguist's claims about which sentences are acceptable constitute predictions about what a speaker population is likely to judge. Often enough, linguists are more rigorous, presenting experimental participants with linguistic stimuli and asking for evaluations—for example, by offering a Likert scale to score each sentence's acceptability.[6]

A similar **coherence judgment task**[7] can be utilized for collecting data for semantic theorizing. As with the acceptability judgment task, data are presented by authors as facts about an expressions, but implicitly this invites readers to make (similar) judgments. Consider a classic case. Given a sentence like (6), speakers infer (7–11), but not (12) or (13).

6 Debates in linguistics have called into question the reliability of trained linguists' judgments, as compared to naive speakers (see Gibson & Fedorenko 2010, 2013 and responses from Culicover & Jackendoff 2010, Sprouse & Almeida 2013). In the end, there is room for multiple methods to advance our knowledge. Evidence suggests that linguists' acceptability judgments converge with naive speakers' judgements more generally (Sprouse et al. 2013).

7 This is more commonly known as a "truth-value judgment task". For reasons that will become clear in §3.3, describing the task in this way is misleading.

(6) Charles stabbed Daniel quickly with a knife.

(7) Charles stabbed Daniel.

(8) Charles stabbed Daniel quickly.

(9) Charles stabbed Daniel with a knife.

(10) Charles stabbed with a knife.

(11) Charles did something with a knife.

(12) Charles stabbed a knife.

(13) Charles jumped Daniel.

Put in a form amenable for acceptability judgments, one constructs contrastive conjunctions by (for example) appending the negation of (7) to the sentence in (6); see the result in (14). This seems to generate incoherence. This unacceptability becomes even clearer if we compare within a minimal pair: (15) seems fine while (14) seems unacceptable:

(14) # Charles stabbed Daniel quickly with a knife, but Charles did not stab Daniel.

(15) Charles stabbed Daniel quickly with a knife, but Charles did not jump Daniel.

Similar pairs could be constructed for conjunctive continuations to (6) using (8–11). The point here is that (14) is unacceptable, despite being perfectly grammatical. This unacceptability is marked with an octothorpe to indicate that the oddity is more like incoherence than ungrammaticality. But really then, this is a proposal about how the unacceptability of these sentences is to be explained, having something to do with the meaning (or use) of the sentences, not their syntactic well-formedness. Such facts require an explanation.

A reformatting of this method closely resembles the method of cases. Presented with a short, fictional **Vignette**, one then probes the intuitions of speakers regarding the expression of interest with a **Prompt**:

> **Vignette**:
> *Charles stabbed Daniel quickly with a knife.*
>
> **Prompt**:
> How well does the following describe the above story?
> *Charles did something with a knife.*

As with Gettier's case above, one is given a (very) brief story, and then asked to evaluate a description of the vignette events. But plainly, the appropriateness of

(11) given (6) tells us nothing about stabbings, anymore than (4) tells us about what dogs can or cannot own. Rather, the data regards a linguistic property, namely what speakers are apt to assert or accept as the result of adverbial modification.

This might come as a surprise to many philosophers, who take the canonical Davidson (1967) to show something about **event structure**. Davidson famously offers a proposal for explaining the linguistic data above that appeals to eventive logical forms. But the sentences in (6–11) themselves do not indicate anything about event structure. Davidson accounts for the inferences English speakers make by first asserting (i) that the meanings of sentences are their truth conditions, and then proposing (ii) that these truth conditions are satisfied by worldly events. The argument in favor of Davidson's proposal, given these hypotheses, is that they can explain the inferential data. But the data presented above do not, absent these assumptions, regard worldly event structure.

Setting aside the merits of Davidson's proposal, Davidson's explanation makes non-trivial assumptions about how the meanings of linguistic expressions are related both to non-linguistic cognition, and the mind-external world. If one takes the verb *stab* as sharing a satisfaction condition with a single concept STAB, which is satisfied by all and only events of stabbing, and takes the meanings that speakers of English understand when they understand expressions that use the verb *stab* to be nothing more than those satisfaction conditions, then the inference data gleaned from (6-15) can regard the content of STAB, and the nature of stabbing events. The same holds for our exemplar case of knowledge. One only gets to claim that cases like Gettier's inform us about our concepts KNOW and the nature of knowledge itself if they take on-board such extra assumptions.

These kinds of extra assumptions highlight the **fundamental problem** with the method of cases. If the method of cases is to provide data beyond that of the acceptability/coherence judgment task, the practitioner has to adopt ancillary **semantic** and/or **ontological** assumptions.[8] One common assumption endorsed by most philosophers is to take sentences to have truth-conditions that worldly objects satisfy (see Bourget & Chalmers (2014) regarding mental content and proper names). Given the situation Gettier's vignette describes, if those sentences are true then the world meets a specifiable collection of conditions. Likewise, the *know*-report about Smith articulates a truth-condition, which either is or is not satisfied by the very situation that satisfies the vignette description. If speakers take the *know*-report to be a felicitous description, then granted the ontological assumptions, this tells us

8 Related points are made within the epistemology literature, and more narrowly targeted at various Knowledge First proposals (Ichikawa & Jenkins 2017). There they highlight the illicit inference that the truth of claims about how agents represent the minds of others entails the truth of claims about the nature of the mental states those represented minds are in.

whether Smith has knowledge. However, Gettier's claims about knowledge itself, and not merely about the use of *know*, require such an extra **ontological** assumption.

More modestly, philosophers make use of the method of cases in doing **conceptual analysis**, whereby speaker judgments count as data regarding non-linguistic cognition. Gettier cases thereby reveal aspects of how we represent other minds, informing us about concepts like KNOW. Moving beyond the linguistic deliverances of acceptability/coherence judgments to make claims about extra-linguistic cognition requires adopting **semantic** assumptions linking linguistic meanings to conceptual contents. As with the ontological assumption, numerous semantic claims linking language to extra linguistic cognition can do the work of licensing inferences about conceptual content from speaker judgments. The simplest of these is to take the meanings of expressions to simply be conceptual contents. The meaning of *know* on this Lockean picture makes the judgments of speakers directly pertinent to the contents of some single concept KNOW. This kind of semantic assumption seems more modest compared to an ontological analog, but it is not trivial. As we argue below, natural language polysemy (§3.3) and the pragmatics of communication (§4.3) tell against this optimistic assumption. If one takes seriously the idea that linguistic meaning is an aspect of a speaker's understanding of expressions, this requires that at least some sort of relationship holds between sentential meanings and the contents of the thoughts constituted by concepts. But the Lockean claim is much more substantial (Vogel 2016).

Irrespective of the plausibility or implausibility of these additional assumptions, in order for the method of cases to offer extra-linguistic data to out-strip the deliverances of the acceptability/coherence judgment task, either additional semantic and/or ontological assumptions are required. This poses the fundamental problem with the method of cases. Since these optimistic assumptions are rarely made explicit, and are universally non-trivial, the use of the method of cases becomes equivocal in their absence. In the absence of a clear articulation of the semantic and ontological assumptions adjoining any use of the method of cases, the results gleaned from speaker judgments will have no clear import for claims about either extra-linguistic cognition or metaphysics.

For the ordinary language philosopher looking to resolve the productive limit problem by appealing to ordinary language usage, these unstated assumptions are important. In the next section we briefly review the kinds of strategies that contemporary ordinary language philosophers propose to address the productive limit problem. The purpose of doing so is two-fold. First, these proposals do not directly address the fundamental problem. This is of course understandable, since this is not their aim. Nonetheless, irrespective of which solution(s) an ordinary language philosopher might deploy, some collection of semantic-ontological assumptions is required. This informs the second purpose of surveying extant ordinary language

philosophy proposals, which is to underscore a mandate common to each proposal, that the method of cases benefits from attending more closely to the cognitive resources utilized in language use. This mandate, we argue, tells against accepting the optimistic assumptions that would address the fundamental problem.

## 3.2 Proposed solutions

### 3.2.1 The experimental solution

The first response to the productive limit problem probes the univocality or veracity of audience judgments/intuitions by simply broadening the audience. Call this the **experimental solution**. One influential version of this response is articulated in Weinberg et al. (2001), who surveyed people with diverse ethnic and socioeconomic backgrounds to see if they judged Gettier-style cases in the way canonically accepted by philosophers. What they found with their surveys was that different groups were differently disposed to count cases of epistemic luck as knowledge. The theoretical upshot is that the very judgments which led philosophers to abandon the orthodox analysis of knowledge are not as stable as philosophers claim.

Further research in this experimental vein has contradicted these initial findings, using similar, and in some cases, identical materials (Machery et al. 2017, Kim & Yuan 2015, Turri 2013). For example in Machery et al. (2017), participants were recruited from 4 distinct global regions and given vignettes describing Gettier cases: "knowledge" cases, and "false belief" cases. The central question posed to participants made use of the attitude verb *know*.[9] The reported finding is "that across quite different cultures and languages people exhibit Gettier intuitions" because (broadly) participants indicated that *know*-reports were not apt descriptions for those vignettes that described cases of epistemic luck (Machery et al. 2017: pg. 651).

On the experimental solution to the productive limit problem, the ordinary language response is to simply expand the method to include ordinary language users. As traditionally utilized, the method is mistakenly relegated to the pedantic corners of analytic philosophy. So, the reply insists that achieving univocal judgments/intuitions to cases requires broadening the pool of respondents to ensure that philosophers' representations are not atypical.

9 For example, after reading a story about a woman buying a fake diamond (the "false beliefs case") participants were asked: "In your view which of the following sentences better describes Emma's situation?":

(1) Emma knows that the stone is a diamond.

(2) Emma feels like she knows that the stone is a diamond, but she doesn't actually know that it is. (Machery et al. 2017: pg. 649)

However, this merely replicates the fundamental problem with the method of cases. Failing to articulate the semantic/ontological assumptions on a given use permits the same broad equivocation between three prima facie distinct objects of study: linguistic expressions, concepts, and mind-independent things. Machery and colleagues claim "that across quite different cultures and languages people exhibit Gettier intuitions" (Machery et al. 2017: pg. 651). But what do "Gettier intuitions" regard? That is, what do these responses provide information about? Machery and colleagues at various points seem to claim that such intuitions are about each of these three objects. Initially, the claimed target of these intuitive judgments is clear, "that Gettier intuitions may be a reflection of an underlying innate and universal *core folk epistemology*" (Machery et al. 2017: pg. 652). Presumably this shared innate system traffics in representations like a concept KNOW, and so Gettier intuitions inform us about such concepts/representations.[10] But earlier in the analysis they state that "Gettier intuitions are judgments indicating that the protagonists in such cases do not have knowledge" (Machery et al. 2017: pg. 651). This seems to be a claim about the nature of knowledge states, not merely about how participants represent others' minds.

These various ways of discussing the findings seem to conflate three prima facie different things: natural language expressions like *know*, the conceptual resources used to represent the mental state of others like a concept KNOW, and the mental states themselves. This conflation seems responsible for assertions like "[o]ther plausible components of core folk epistemology are evidential markers" (Machery et al. 2017: pg. 652). As they make clear, "core folk epistemology" is a cognitive system used to think about mental states like knowledge and belief. But evidential markers (found in languages like Korean and Turkish) are morpho-syntactic features of linguistic expressions that grammatically encode evidential sources. Whatever one might want to say about evidential markers and their relation to the concepts we use to think about the mental states of others, it is difficult to understand such markers as parts of concepts.

Nonetheless, Machery et al. (2017) represents a more sophisticated version of this first response to the productive limit problem, which makes use of more sophisticated methodologies from the social and cognitive sciences, introducing careful controls, balanced interventions, and other means of offering more targeted judgments/intuitions from an audience.[11] Whatever one might want to say about

10 Indeed, they say exactly this: "the 'method of cases' which uses intuitions about hypothetical cases as evidence in evaluating analysis of philosophically important concepts, like the concept of knowledge…" (Machery et al. 2017: pg. 646).

11 For example, Turri (2013) illustrates that a structural repackaging of Gettier cases biases respondents' judgments about cases of epistemic luck. In one study, two groups were asked to respond to the same Gettier cases, with Turri manipulating the temporal presentation of the text. For one group, the

the merits of contemporary experimental philosophy, this kind of response to the productive limit problem treats shortcomings with the method of cases to be a failure of application, not a failure in the method. But such a response ignores the fundamental problem, which is simply recreated by incorporating a larger population of speaker judgments.

### 3.2.2 The natural solution

A second kind of response to the productive limit problem identifies the shortcomings of the method of cases with its artificiality. Here, the focus is on the **naturalness** of the vignettes being evaluated. The kinds of vignettes philosophers develop are, on this response, overly technical, abstract, and peculiar. More ordinary cases, ones that better reflect the ordinary use of language, are more likely to recruit participants' competence directly. Call this the **natural solution**.

An example of this strategy is the work of DeRose (1992), who utilizes the method of cases to argue for epistemic contextualism.[12] Here the vignettes reflect common activities for modern people where realistic concerns are compared across cases. DeRose contrasts two situations where the evidence proffered as justification for the protagonist's claim to knowledge remains fixed, along with the mentioned possibility of error, but the stakes for being mistaken differ. The protagonist in these vignettes is concerned with the closing time of a financial office at which they hope to make a deposit, claiming that "I know [the bank to] be open. I was just there two weeks ago on Saturday" (DeRose 1992: pg. 913). In the first telling, the deposit is not terribly important, while in the second telling the deposit is much more so—failing to clear before the weekend would have dire consequences. The lesson proffered from these naturally contrasting cases seems to be that whether a belief counts as justified depends contextually on what is at stake, modulating whether one does (or does not) have the professed knowledge.

The cases utilized by DeRose (1992) stand in contrast with the peculiarity of Gettier's case (and the fantastical oddity of others in the epistemology literature, e.g., Lehrer & Cohen 1983). Gettier's Smith forms a fairly odd belief, using a baroque description of a man he can more directly represent. Such a belief is decidedly unnatural, but central for generating the epistemic luck in Gettier's case, since the

Gettier cases were presented in three distinct stages, dividing the text into parts corresponding to the structural features common to Gettier cases. A second group were presented the same vignettes in a single stage. The manipulation to the temporal presentation of vignettes had the effect of aligning participants' judgments with the judgments of philosophers.

12 The epistemic contextualist denies that epistemic terms have context-invariant meanings. One such account treats words like *know* as indexicals like *here*. Just as the significance of a particular use of *here* will change depending on the location of its use, the significance of *know* varies from context to context (Cohen 1988).

coin-based description applies equally well to both Smith and his competition. This feature of Gettier cases gets exacerbated in counter-examples to the increasingly complex proposals offered to close the flood gates Gettier opened. The natural solution implicit in DeRose (1992) to the productive limit problem identifies the problem with this extravagance.

As with the previous ordinary language response, the natural solution takes the method of cases to be fundamentally reliable.[13] Here the concern is not with whether careful experimental controls are implemented, but rather with the **naturalness** of the vignette's content. Presumably, the artificial character of standard vignettes privileges particular interpretations, and fails to capture to breadth of a participant's competence. The more natural the content of the described case is, the more likely it is to reflect the robust competence of respondents. But whether the results target knowledge, or epistemic cognition, remains unclear. That speaker responses to *know*-reports vary as vignette contents shift what is at stake does not settle whether "contextually determined variation on how good an epistemic position one must be in to count as knowing" regards knowledge states, or our epistemic representations (DeRose 1992: pg. 614). For speaker judgments to indicate that knowledge is context dependent, further (ontological) assumptions are needed.

### 3.2.3 The expert solution

A third response is, in a sense, to deny the productive limit problem by granting that ordinary language users will routinely diverge in their treatment of various cases. But for this reason we should attend only to the intuitions of experts in using the method of cases. The most recent iteration of this **expert solution** stems from work on **conceptual engineering**. This response takes a central task of philosophers to be "the project of assessing and developing improvements of our representational devices" (Cappelen 2020: pg. 133). Given this task, the method of cases can be utilized to both describe how we represent the world, and inform how we ought to represent it. The normativity of the engineer's mandate is derived from a particular goal or purpose. As such, the productive limit problem reflects subtle disagreements about the purpose of inquiry, not a problem with the method of cases.

Clark & Chalmers (1998) is offered as a paradigm of this improvement strategy (Cappelen 2020: pg. 138). Their central claim is that core mental states, like beliefs, contain mind-external things as constituents. Central to their argument is the method of cases (Vogel 2018). Their vignette describes an individual that is unable to

13 The experimental and natural solutions are not mutually exclusive—both modifications may be needed to respond to the productive limit problem. In fact, Hansen & Chemla (2013) seems to be an instance of exactly this, soliciting judgments from a multitude of participants about natural seeming cases where the stakes vary in natural seeming ways.

construct long-term memories—viz. Otto—who stores information in a notebook as a substitute. The case is offered to show that Otto's beliefs are not merely informed by the notebook, but that the notebook is a constituent of his beliefs—literally a handheld part of his mind. They are explicit about their project, that it is revisionist:

> Some will resist this [revisionist] conclusion. An opponent might put her foot down and insist that as she uses the **term** 'belief', or perhaps even according to standard usage, Otto simply does not qualify as believing... We do not intend to debate what is standard usage; our broader point is [about how] the **notion** of belief ought to be used... By using the 'belief' notion in a wide way, it picks out something more akin to a **natural kind**. The notion becomes deeper and more unified, and is more useful in explanation (Clark & Chalmers 1998: pg. 14, our emphasis).

The normative force of the revision rests on a particular goal—these revisions are **purpose-driven**. For Clark and Chalmers this goal is a philosophically pedestrian one, offering a more parsimonious explanation. But, for those conceptual engineers proposing modifications in other domains, the purpose might be otherwise. For example, a goal of moral-concept engineering might be that the use of certain concepts (and not others) lends to more just outcomes (Haslanger 2000).

But, this expression of the proposal, and likewise the more general defense of the conceptual engineering response in Cappelen (2020), betrays the fundamental problem with the method of cases, mirroring the confusion highlighted in discussing Machery et al. (2017) regarding the relationship between linguistic expressions, concepts, and the things they seem to be about. Clark and Chalmers claim their proposed revision offers a more useful explanation. But what precisely is being explained? Is the target of explanation the content of a concept (a "notion"), the meaning of a linguistic expression ("term") like *belief*, or the nature of mental states (as a "natural kind")? These three are distinct natural objects, and to the degree Clark and Chalmers are offering a coherent proposal, one wants to know what their proposal is about such that the method of cases is informative about that thing.

This confusion is not a mere incidental slip of the proverbial tongue, one that results from a lack of explicit attention to the details of philosophical methodology. Cappelen (2020) describes the general conceptual engineering strategy in equally equivocal ways, which itself is a strategy that regards philosophical methods. In Cappelen's articulation of the conceptual engineer's "master argument" the main purpose of philosophy according to the argument seems quite clearly to regard the meanings of linguistic expressions (viz. "lexical items"/"terms"): "no matter what topic a philosopher is concerned with, they should assess and ameliorate the meanings of central terms" (Cappelen 2020: pg. 135). Yet, throughout the

defense of this conclusion we see a shift in the scope of the argument. At one point Cappelen replies to an objection by indicating that a "purely descriptive project" is not possible. The reason for this impossibility is that the main purpose of philosophy, as the master argument outlines, can in principle be the object of engineered revision: "all the **concepts** involved in describing conceptual engineering should themselves be subject to critical assessment" (Cappelen 2020: pg. 147, our emphasis). If the target specified in the conceptual engineer's methodological mandate are linguistic expressions and their meanings, then a further assumption is required for directing that mandate at the content of the **concepts** those expressions recruit. This kind of equivocation appears elsewhere in detailing the master argument, both in discussing meaning change, and denotation (Cappelen 2020: pg. 132, 136)

We see here the re-creation of the fundamental problem most strikingly, in part because of the explicit attempt to articulate the conceptual engineering project without making any semantic/ontological assumptions (Cappelen 2020: pg. 136). The method of cases, in order to yield verdicts about the to-be engineered conceptual contents, must make some assumption about how linguistic meaning relates to extra-linguistic cognition. But the uncontroversial claim that meanings are somehow related to extra-linguistic cognition is not enough to justify the use of the method of cases for the aims of conceptual engineering.

### 3.2.4 The ecological solution

A fourth response to the productive limit problem focuses not on the language used in the method of cases, but regards the ecological validity of that use. The worry is not that the vignettes utilized in the method are unnatural, but that the task itself is. The solution is to bring the task into a more natural "ordinary" environment. Call this the **ecological solution**.

Hansen (2020) exemplifies a particularly sophisticated version of this kind of reply. The motivation here is inspired by work in other experimental domains that highlights how participants' judgments (about uncontroversially clear situations) can be influenced by a number of factors: experimental settings, participants' motivations within the task, awareness that the task is an experimental one, and others. For example, Baron et al. (1996) studied how participants utilized the opinions of peers on a visual task. They found that participants, when given more time to produce a perceptual judgment, are less influenced by the judgments of others when the task is presented as being consequential. In their task, when participants were told that their judgments would bear on whether a new policing procedure would be implemented in the near future, participants were less swayed by peer opinions, as compared to

controls.[14] This kind of finding bears on the use of the method of cases either from the armchair, or in a sterile experimental context. A participant's case judgment no doubt seems to them to be utterly inconsequential, given that they are evaluating a fictional case read in complete isolation from any practical concerns. The idea then is that merely presenting a participant with a vignette leaves out an important input to the cognitive processes engaged in representing situations—namely, motivation. This kind of intervention, that borrows from extant methods in the behavioral and cognitive sciences, seems more useful than merely amplifying the method of cases by getting judgments from ordinary language users. As we highlight in §§4 and 5, similar concerns beset the use of language in methodologies in the cognitive and social sciences, which admit to various pragmatic confounds.

But even with this nuanced reply, the fundamental problem with the method of cases remains. Suppose we implement a variety of interventions that more aptly capture the ecological variability that agents find themselves in when making epistemic judgments. And we develop more natural vignettes, applied to a diverse population of participants, implementing the experimental, natural, and ecological solutions to the productive limit problem. And suppose, as some evidence currently suggests (Hansen & Chemla 2013, Francis et al. 2019, Dinges & Zakkou 2021, Grindrod et al. 2019), that the results fall in line with what epistemic contextualism would predict: we find that subjects' reported judgments about epistemic situations predictably vary along multiple context-variant dimensions. Which of the following broadly contextualist views would this data support:

- the concept KNOW is context-sensitive (Glanzberg 2011)[15]
- the word *know* contains or generates context-sensitive elements (Cohen 1988, 1999, DeRose 1995, Lewis 1996)
- mental-states of knowing are sensitive to contexts (Annis 1978, Williams 2001)

14 The findings here are a bit more nuanced. Four conditions were tested modulating both the significance and the difficulty of the perceptual judgment task, where a temporal constraint served as a proxy for difficulty. After all, making a perceptual judgment with less time is more difficult. The finding was that, the influence of raising the importance of the task only manifested in the less difficult condition, where participants had more time to make their judgment. Intuitively, participants were more swayed by peers if they either took the task to be unimportant, or if they took their judgment to be less reliable given the difficulty of the task. But the sway of their peers diminished if the task was both important and less difficult.

15 Glanzberg (2011) does not directly discuss *know* or its correlative concepts. But the view described therein takes "root" word meanings to contain extra-linguistic conceptual content that "packages" the resources of extra-linguistic cognition in ways that are (broadly) context sensitive. See also Glanzberg (2014, 2018).

- context determines which one of a variety of homophonous linguistic expressions *know* is being used in an utterance (Williams 2002)
- context influences, without determining, which of a variety of KNOW concepts gets retrieved upon executing the instructional meaning of a single lexical entry for *know*

The method of cases is utilized under the broad assumption that the verdicts of participants to prompts (and their related vignettes) in a natural language are indicative of something, but more care needs to be paid to attending to what these responses are indicative of.

### 3.3 Why the fundamental problem matters for ordinary language philosophy

A philosopher at this point might be unmoved by the fundamental problem for the method of cases. Granting the problem, they might simply make explicit their semantic and ontological assumptions, avoiding the worrying equivocation.[16] If one adopts the semantic-ontological assumption to "think, in common with many philosophers, of meanings as at least intensions, that is, functions from points of evaluation to extensions" then we find a resolution of the fundamental problem (Cappelen 2020: pg. 136). If we take Cappelen's assessment of the field as apt[17], then many take the following assumptions as a prerequisite for engaging in the method of cases:

For any expression *e*:

(EC) the meaning of *e* determines the concept *e* expresses

(ET) the meaning of *e* determines the truth-conditions *e* expresses

The method of cases, given these assumptions, addresses linguistic meaning, conceptual content, and ontology in one fell swoop. If the connection between the meaning of *know* and a concept KNOW is that they share an intension that determines an extension, then talk of expression meanings and conceptual contents is straight forwardly interchangeable—contents and meanings are both intensions (viz. functions/satisfaction-conditions) that connect them to a worldly domain. Such assumptions may seem natural to many philosophers. Indeed, much of the confusing equivocation we've highlighted in Machery et al. (2017) and Cappelen (2020) can be rendered fairly consistent if one adopts such a truth-conditional proposal about the

16 Our thanks to an anonymous reviewer for raising this concern.

17 Others corroborate this sociological claim, that most philosophers are externalists about meaning and content (Gertler 2012).

relationship between expressions, concepts, and their purported extensions. If one takes "the concept expressed in English by *knows*" to have a determinate extension that indicates "which Gettier cases are and are not instances of knowledge" then participants' responses to Gettier vignettes can directly inform us about the meaning of *know*, the content of KNOW and the nature of knowledge states (Machery et al. 2017: pg. 652).

But the assumptions in (EC) and (ET) admit to immediate counterexamples given indexicals/demonstratives (e.g., *I*, *this*, *today*, *last month*; see Strawson 1950; see Lewis 1980 for an illuminating discussion) and gradable adjectives (e.g., *blue*, *tall*, *expensive*; see Kennedy 1999). The sentence *This is expensive* offers no determinate truth condition, in the absence of a context to satisfy features that vary along at least two dimensions.

In response one might embrace a more elaborate version of our semantic and ontological assumptions that takes contextual variation into account.

For any utterance of expression $e$, in context $c$:

(ECC) the meaning of $e$ determines the concept $e$ expresses in $c$

(ETC) the meaning of $e$ determines the truth-conditions $e$ expresses in $c$

Assumptions of this sort seem to fit quite well with what many ordinary language philosophers want to say about the context sensitive nature of knowledge. Despite this, these are bold and contentious assumptions about linguistic meaning (Pietroski 2018), conceptual content (Quilty-Dunn 2021) and metaphysics (Vogel 2020). We'll suggest one source of concern here, regarding natural language **polysemy**. The purpose here is not to show that theses like (ECC) and (ETC) are false, but to show that, as solutions to the fundamental problem, they come at a cost, and in particular at a cost the ordinary language philosopher may not want to pay.

Natural language polysemy suggests that, even constrained to a single context, an utterance of an expression will not determine a unique concept/satisfaction condition. Take for example the expressions

(16) Seattle is friendly, and it is rainy.

(17) Seattle is friendly.

(18) Seattle is rainy.

One might make the claim in (16), using the polysemous expression *Seattle*, while discussing the relationship between weather and human dispositions. Any English speaker would be apt to infer both (17) and (18) from (16). Such data goes unexplained by our assumptions in (ECC) and (ETC). The single use of *Seattle* licenses the recruitment of two distinct concepts (or the correlative satisfaction conditions):

one for thinking about a population, and another for thinking about a geographic region. Since the anaphoric *it* must borrow its truth conditional contribution to (16) from that of the antecedent *Seattle*, whatever truth-condition (16) expresses in this single context, only one of these satisfaction conditions can be recruited according to (ECC)/(ETC). Thus, whatever single truth-condition a use of (16) would determine in context (either regarding a region or its inhabitant), that truth condition, if satisfied, would fail to entail either (17) or (18). Thus, however speakers infer (17) and (18) from (16), it cannot be a matter of logical entailment from the truth-condition that is determined by a use of (16) in a single context. This runs counter to the predication of (ECC)/(ETC), which accounts for such inferences as logical entailment.

There are various strategies one might invoke to respond to the problem polysemy poses (see Falkum & Vicente (2015), Gotham (2017), Vicente (2021), Gotham (2022) for helpful discussions). A natural way of resolving the concern that polysemy poses for hypothesis like (ECC) and (ETC), is to build the flexibility of such expressions into lexical entries for them, treating them like indexicals. In fact, this seems to fit well with much ordinary language philosophy work (discussed above) which is taken to support context-sensitive proposals regarding knowledge and knowledge-attribution. Lexical entries for polysemous expressions would then (somehow) contain variable elements that get filled in by different contexts in different ways, just as the expression *I* is satisfied by different individuals in different contexts depending on who is speaking (see Rothschild & Segal 2009 for such a proposal using color terms as a paradigm).

There are however, general concerns about such a strategy. For one, this requires adopting a permissive notion of context, one that is in tension with the motivations for distinguishing context-sensitive expressions like indexicals from other broadly pragmatic phenomena (Kaplan 1989). These concerns are amplified if one takes this kind of indexical approach as a claim about linguistic cognition. Plausibly, the human language faculty would have developed a system to pay special attention to the speakers, times, and locations of utterances—and thus developed a system that directly exploits this information in constructing lexical items like *I*, *here*, and *now*—since all utterance contexts supply this information. It is far less plausible that this system would be sensitive to the much wider range of information required if every lexical entry admits to indexical complexity. Second, given the prevalence of polysemy, which represents the most common kind of ambiguity in natural languages (Apresjan 1974, Ostler & Atkins 1991, Pustejovsky 1995), this proposal would require that not only is most of the lexicon populated by indexicals, but those indexicals can shift reference mid-sentence. This not only presents substantive hurdles for language acquisition (Husband et al. 2016) but would also fail to explain why children distinctively struggle to learn paradigmatic indexical expressions like *today*, *tomorrow*, and *yesterday* (Grant & Suddendorf 2011, Weist et al. 1991).

Alternatively, a historically favored solution to problems like those posed by polysemy is to reject the relevant speaker data, by claiming that speakers are simply mistaken in their judgments, either systematically or within the relevant domain. This reply is grounded in a commitment to taking languages to be E-languages, mind-external things that speakers can represent in better of worse ways.[18] Offering a theory of meaning on such views requires modeling the mind-external thing that speakers learn (sometimes/often? poorly and incompletely). We suggested at the outset some reasons to be wary of such a view of languages.

But setting aside the merits of these kinds of responses, the ordinary language philosopher ought not embrace them, as doing so would run counter to their commitments pertaining to the productive limit problem. What we see in the various ordinary language philosophers solutions to the productive limit problem (barring the expert solution) is that the method of cases, as traditionally practiced, fails to achieve its purported end, of either revealing the content of concepts, or the nature of the things in their extension. The diagnosis for this failure is that such traditional use fails to reflect how humans conceptualize the domain of interest—with our exemplar case being knowledge. The experimental solution is motivated by the worry that the intuitions/judgments of philosophers might be aberrant, failing to reflect how humans generally conceptualize knowledge. The natural solution charges the method of cases with utilizing stimuli that curtails a more comprehensive recruitment of the conceptual resources utilized in attributing knowledge. The ecological solution presses on that artificiality of the testing environment, which introduces biases that impact how one processes and responds to stimuli. What unifies these responses is that the method of cases traditionally fails to sufficiently attend to the participant's cognition, an thereby fails to reflect the underlying conceptual systems used to represent the relevant domain—in the case of knowledge, the cognitive systems used to represent the minds of other agents. The ordinary language philosopher's hope is that appealing to facets of language use, in one way or another, corrects this failing.

Taken together, the ordinary language philosopher's response to the productive limit problem is a mandate to (i) more closely attend to how human cognition operates, by (ii) appealing to the facets (and richness) of language use that the method of cases traditionally eschews. The upshot of this mandate is that commensurate attention must be paid to how human cognition operates with respect to language itself. If the rich and varied nature of language use is to address the productive limit problem, that richness is tied to the diverse interpretations speakers apply to identical linguistic inputs. If a word like *Seattle* is capable of soliciting multiple satisfaction conditions, one needs to know how the meaning of *Seattle* makes that possible.

18 See Quine (1960), Davidson (1973b), Lewis (1975), Dummett (1978), Devitt (2006). For this kind of global response to vagueness see Williamson (1994). For a contemporary response to polysemy of this kind see Burge (2003), Kennedy & Stanley (2009).

One needs to know, what is it that children learn when they acquire the meaning of such expressions. Both the E-language inspired denial of the polysemy data and the indexicalist account render the task of addressing that acquisition question either perplexing or intractable. Further, since a speaker's declarative uses of (18) can offer an imperative pertinent to packing for a trip (Searle 1975), one needs to know what the sentence itself provides to a cognizing audience such that pragmatic processing can get them to arrive at a thought that more closely resembles one recruited by an imperative: *Pack an umbrella*. Put simply, the ordinary use of language involves speakers utilizing procedures to generate and comprehend the utterances of other speakers. This tells against accepting the view that the study of language is the study of E-languages, but rather concerns the procedures speakers implement, and children acquire, during linguistic cognition.

Our purpose is not to argue against externalist theses like (ECC) and (ETC).[19] Rather, the concern is that the fundamental problem with the method of cases is not easily resolved for the contemporary ordinary language philosopher. Resolving the fundamental problem by appealing to claims like (ECC) and (ETC) burdens the ordinary language philosopher with the task of addressing problems like those that polysemy poses for such claims. The most straightforward ways of addressing the fundamental problem for the method of cases are ill-suited to the ordinary language philosopher's mandate in response to the productive limit problem.

This bears on how an ordinary language philosopher approaches the empirical work reviewed in the remaining sections. Given the use of natural language stimuli in the methods discussed below, it is worth attending to the operative notion of language implicit in such use. Taking languages to be procedures for generating meaningful expressions informs how the results of empirical inquiry can inform the questions philosophers care about.

## 4 Studying the concept

Similar problems to those pointed out for speakers' judgments in (§3.1) also apply to methods that rely on speakers' behaviors more generally (be they judgments of acceptability and coherence, interpretations of others' utterances or even productions of one's own utterances). Given that the objects of study are mental objects, it can be hard to probe them directly, as opposed to examining them through their relationship to other cognitive processes or mind-external things. In the specific case of KNOW, empirical work with verbal populations largely relies on the interpretation of sentences or utterances involving *know*. For example, we might ask a participant whether someone really *knows* something or whether they merely *believe* it, on the

19 For efforts in that vein see Chomsky (2000), Pietroski (2005), Valente (2019), Vogel (2016).

assumption that *know* has a deep relationship to KNOW, and *believe* to some other concept such as BELIEVE. But this useful empirical shortcut can lead us astray if we fail to recognize that language is supported by its own mental representations and cognitive mechanisms, which could differ from those utilized in extra-linguistic mental state attribution. In using linguistic stimuli, we are thereby probing the mental processes that underlie both language and mental state attribution simultaneously. As such we should take pains to be explicit about how these two distinct cognitive systems interact with each other, especially if the goal is to make inferences about one of them in particular. In this section, we will provide a window into work from developmental psychology targeting the development of the ability to attribute knowledge. Through this window, we will discuss what this work can tell us about the concept, and the importance of linking assumptions about the relationship between concepts, linguistic items and the empirical tools we use to test them.

### 4.1 What is KNOW?

The field has not yet arrived at an explicit formal proposal which cleanly distinguishes KNOW from other concepts by articulating the different kinds of mental objects which underlie knowledge attribution. However, Phillips et al. (2021) have recently taken admirable steps towards defining the object of study, by articulating four essential characteristics that knowledge representations should have.

(19) Features of knowledge representations:
   a. they are necessarily veridical;[20]
   b. they are distinguishable from true belief representations;
   c. they allow for attributing both altercentric and egocentric ignorance;
   d. they are not tied to a specific modality.

These four features give us a concrete definition which matches our intuitive judgments of what distinguishes knowledge attributions from other closely related kinds. First, knowledge attributions are distinguished from belief attributions because the former's content must match reality—and are thereby veridical—while the latter's need not. Second, attributions of knowledge seem to require something more to them than attributions of true belief, as Gettier cases suggest. Third, attributions of knowledge must allow us to distinguish what we know from what another knows, even when we lack their known information. And finally, the knowledge that is

20 Phillips and colleagues use the term "factive", a borrowing from linguistics which was originally coined to label a class of presupposition-triggering predicates, including *know*, inspired by the syntactic properties of those predicates (Kiparsky & Kiparsky 1970). Given that factivity has a different meaning in the linguistic literature (see §5) we prefer to use the term "veridical" to avoid confusion.

attributed can be updated via any relevant (perceptual) source. In the next section, we will briefly review the evidence that exists for representations that meet these four criteria in infants and young children.

The point of reviewing the findings in the next section is two-fold: to offer a broad indication of the methodologies used in investigating epistemic cognition, but also (as we argue) to underscore the confounds that **pragmatics** introduces to these methodologies that crucially rely on the interpretation of linguistic stimuli (whether or not the desired object of study is language).

## 4.2 Evidence for knowledge representations

### 4.2.1 The veridicality of knowledge representations

First let's turn to the necessary veridicality of knowledge representations. There is a significant amount of work examining the development of the distinction between attributions of mental content that are congruent with reality vs. incongruent with it, given the historical interest in false belief cases.

Foundational studies in the Theory of Mind literature were focused on the attribution of false beliefs as a test case for metarepresentational abilities (following Premack & Woodruff 1978 and associated commentaries). The ability to attribute false beliefs is informative about the ability to metarepresent because attributing a mental content which is not veridical requires another representational layer to delimit or quarantine the mental content being attributed from other representations (i.e., those of reality). But, important for our purposes here, measuring the ability to attribute such non-veridical mental contents also requires control conditions where participants are asked to attribute veridical mental contents.

One traditional task for testing false belief attribution is the "change of location" task. In reality-congruent conditions, an agent thinks that an object is in location A because they saw it there. In reality-incongruent conditions, the object is displaced to location B when the agent is absent, making their belief inconsistent with the current state of reality. In these **elicited-response tasks**, verbal children are asked questions like *Where will the agent look for the object?* The well-replicated finding is that four-year-olds respond in an adult-like way to the agent's belief regardless of the actual facts of the matter, while younger children fail to respond to anything but the actual state of affairs (Wellman et al. 2001, Perner et al. 1987). Taking this perspective, we could assume that children under 4 years are not able to distinguish veridical attributions from non-veridical ones, perhaps because they lack the capacity for building non-veridical ones.

But the orthodox view from these elicited-response paradigms has been challenged by those who argue that the results are due to performance errors related to the

response prompt. One salient possibility is that because responding to the linguistic form of the prompt recruits the language faculty, we must distinguish between the roles that the language faculty and extra-linguistic conceptual competence play in children's responses (see Dudley 2018 for discussion). As we argue in §4.3 this includes accounting for the pragmatic competence of participants, a concern that aligns with the ordinary language philosopher's contention about the informativeness of language use in inquiry. **Spontaneous-response tasks** try to remove this potential source of errors by asking whether young participants can differentiate reality-congruent vs. reality-incongruent attributions when they are not provided with an explicit linguistic prompt. Instead, these tasks examine non-linguistic behavior via diverse paradigms and dependent measures (Onishi & Baillargeon 2005, Clements & Perner 1994, Buttelmann et al. 2009, Moll et al. 2016, Southgate & Vernetti 2014, Barrett et al. 2013). Unlike elicited-response paradigms, this body of research suggests that—even as young as 6 months—children can represent mental contents which match reality as distinct from those which do not match reality.

What should we make of these conflicting bodies of results? On the one hand, we could explain away early successes in spontaneous-response tasks (partly motivated by the relative dearth of replications thus far, see Dörrenberg et al. 2018 but also Baillargeon et al. 2018) and take the position that the ability to distinguish veridical mental contents from non-veridical ones emerges relatively late in development. This position also leaves some room for different positions on knowledge: are younger humans unable to differentiate representations on the basis of their veridicality because they can only entertain knowledge attributions, or rather because the distinction between knowledge and belief has not yet emerged for them?

On the other hand, it could be that the ability to represent veridicality vs. non-veridicality emerges early in ontogeny and failures in elicited-response tasks are explained by orthogonal competences, which are developing independently but simultaneously (Leslie et al. 2004, Carlson & Moses 2001, Siegal & Beattie 1991, among others). In particular, if language comprehension is required by elicited-response tasks but not by spontaneous-response tasks, we can ask how children's developing linguistic competence contributes to the differences in their behavior within these two paradigms. See §4.3 for further discussion of this possibility.

However conclusive one takes this literature to be, it is important to notice that it only tells us about **reality-congruence** (Nagel 2017): we can take the findings to reflect the ability to attribute veridical vs. non-veridical mental content. Most Theory of Mind studies only license conclusions which distinguish veridical attributions (reality-congruent ones) from others (reality-incongruent ones), but both knowledge and true belief attributions (when apt) are veridical.

#### 4.2.2 Attributions of knowledge vs. true belief

Studies that try to differentiate between attributions of knowledge and attributions of true belief need to operationalize the difference in some way. One method targets justification. Such studies suggest that children understand that perceptual access can justify knowing something, but they may not distinguish between better or worse sources of justification when tested via verbal reports. Using non-linguistic stimuli, Kaminski et al. (2008) used distinctions in perceptual access to create situations compatible with attributions of knowledge ("known lift" condition), ignorance ("unknown shift"), or true belief ("unknown lift"). They found that children differentiated between these three cases. In contrast, studies using linguistic stimuli find more mixed patterns of success and failure. Children may report having *seen* things they were only *told* or *guessed* correctly (Woolley & Bruell 1996). If given conflicting reports about the location of a toy, such as *Nico guesses that it is in the box* and *Kevin knows that it is in the bag*, they will not preferentially search in the location reported in the *know* sentence (Moore et al. 1989). When asked if they *knew*, *remembered* or *guessed* some proposition that they endorse, they assent to all three characterizations even though they are seemingly incompatible (Johnson & Wellman 1980). However, as researchers behind these findings point out, we cannot know whether children's responses come from (i) their inability to distinguish justified mental states from unjustified ones or (ii) their developing understanding of the meanings of attitude verbs.

#### 4.2.3 The relation between knowledge attributions and sources of information

While knowledge should be modality-general and can be attributed on the basis of many information sources (e.g., vision, touch, inference), attributions of perception are presumably modality-specific, involving various kinds of representations linked to different perceptual sources (e.g., visual evidence distinguished from auditory evidence). What evidence do we have for children's ability to attribute knowledge as modality-general?

Behavioral experiments suggest that young children and infants understand the role of perceptual access in acquiring knowledge. Children themselves acquire information through perceptual access (Call & Carpenter 2001, Gopnik & Graf 1988). They are sensitive to what others are able to perceive, and use that information to navigate interactions (O'Neill 1996). They attribute knowledge to others who have had perceptual access to the relevant information, but not those without perceptual access (Pillow 1989, Sodian et al. 2006). And children understand that perceptual access is necessary to attribute knowledge but not to attribute imaginings (Woolley & Wellman 1993).

There then seems to be an argument for the modality-generality of children's knowledge attributions, especially since they cannot always report which modality was the source of the information that they have gathered (Gopnik & Graf 1988, O'Neill & Gopnik 1991). Perhaps they do not even understand the link between perceptual modalities and the kinds of information that each modality is best suited to provide (e.g., touch is not usually informative about color, see O'Neill et al. 1992, O'Neill & Chong 2001). Or perhaps they do not form strong representations of information sources, because they cannot always report how they acquired information (Papafragou et al. 2007b) or they struggle to make correct judgments about how likely others are to have that information (Taylor et al. 1994)—although this is an error that adults can also make (see Thomas & Jacoby 2013 for discussion).

#### 4.2.4 Attributions of knowledge and ignorance

An important aspect of knowledge representations is that they also allow for attributions of ignorance, which require an asymmetry between the self and another. The ability to attribute ignorance to oneself seems to be present early in development and throughout, which may not be surprising if we need to find knowledgeable others who can teach us about our surroundings. Infants as young as 12 months recognize when others have more information than themselves, and they use this information to decide who to learn from (Kovács et al. 2014, Begus & Southgate 2012). Later, two-year-olds use non-verbal communication to admit that they lack information when they are questioned (Harris et al. 2017a) and three-year-olds take their own (lack of) information into account when communicating with others (Kim et al. 2016). More recently, Kovács et al. (2021) have shown that 15-month-olds track others' knowledge even when the infants themselves cannot know what it is that the other knows.

In contrast, Wimmer et al. (1988) suggested that the capacity to attribute ignorance to others was not in place before 4 years. However, later work by Pratt & Bryant (1990) removed methodological confounds and found that three-year-olds can make such attributions. This capacity to separate one's own knowledge from others', however, may not be perfect. Studies that explore "the curse of knowledge" suggest that three-year-olds are likely to attribute the knowledge that they possess to another (Birch & Bloom 2003), but the explanation of these errors may not require positing a developmental change (Birch & Bloom 2007).

### 4.3 Language-based methods and pragmatic confounds

It is useful to dwell on the role of language in this work on the understanding of knowledge attributions. In this section, we give some examples to highlight

how unstated assumptions about the relationship between language and thought get smuggled into interpretations of findings. Elicited-response tasks employ language in order to first establish the experimental conditions and then later to give participants prompts to respond to. Even though the desired object of study is conceptual understanding, few methods allow us to target it directly without presenting linguistic stimuli to a participant. This carries with it assumptions about how participants are interpreting the utilized linguistic expressions, so we cannot ignore the role that linguistic understanding plays in the results that we get.

In some cases the language employed in a task may just be too complex (e.g., syntactically) for children at the age being testing, leaving too many inferential degrees of freedom. Take for example, a study by Wimmer et al. (1988). Pairs of children played a game where one child (Anne) was able to look inside a box to see the contents and then the other child (Bill) was asked *Does Anne know what is in the box or does she not know?* Wimmer et al. found that the majority of four-year-olds would correctly attribute knowledge to the other child in these instances (by responding "yes"), while three-year-olds would not attribute knowledge to the other (by responding "no"). But there are several reasons why a conclusion about knowledge attribution from a simple yes/no response in this case is too hasty. First, children were asked to respond with a single polar response particle like *yes* to a disjunction of polar questions (like $p? \vee \neg p?$). But what does an affirmative response indicate? These children are not likely assenting to (or denying) the disjunctive proposition. And if they were, nothing about their conceptual competence is revealed by the task. They are almost certainly assenting to or rejecting one of the disjuncts. But why assume that the first is targeted by an affirmative response and not the last. Second, negative polar questions give rise to implicatures (Romero & Han 2004) and single-particle responses can be ambiguous without further clarification:

(20) George isn't from the UK?
  a. No (... he's not).
  b. No (... he is).

This leads to a third concern, as we detail in §5, that some children do not have an adult-like understanding of the semantics of *know* at age 3 (Moore & Davidge 1989, Dudley et al. 2015). If they are still coming to understand the meaning of *know*, then their response to a question deploying this lexical item may not recruit their ability to attribute knowledge, regardless of the relationship between words and concepts.[21] Fourth, there is also the orthogonal methodological consideration

21 Consider your own response to the question *Does Anne gorp what is in the box?* and what that response has to do with your ability to attribute knowledge to others. This kind of context has been used to simulate how adults with full mindreading competence will respond to different kinds of linguistic stimuli (Gillette et al. 1999).

that young children's responses to polar questions do not always reflect deeper understanding, but simply betray a response bias (Mehrani 2011). The force of these concerns is supported by the fact that a follow-up study by Pratt & Bryant (1990) using a different test question found a different pattern of results.

With just these few points on the table, we start to see that there are many factors to consider for just one question in one study. But this case is instructive. If the task is intended to investigate the conceptual repertoire of developing children with respect to their ability to represent states of knowledge, soliciting responses to linguistic items invokes the cognitive systems recruited in representing states of knowledge but also the systems for comprehending language. Assuming that children's verbal response to such queries directly reflects their conceptual competence betrays a commitment to the same optimistic assumptions implicit in the standard use of the method of cases. Doing so assumes an (overly) simple relation between uses of *know* and how the concept KNOW is deployed in representing the mental states of others.

Even when the linguistic prompt is less complex, hasty conclusions are inadvisable. In some cases, the linguistic prompt is ambiguous, and children's developing pragmatic abilities may lead them to take a different interpretation from the one that the experimenter intends. Take the traditional elicited-response false belief tasks discussed above. Some of these tasks require children to evaluate assertions like (21b) in contexts like (21a).

(21) a. **Scenario**: Danny has already made it home for the day, but Andreas wrongly believes that he is still at work.

b. **Test sentence**: Andreas thinks that Danny is at work.

The typical finding is that four-year-olds (and older participants) will accept this assertion, which is the correct behavior given that it accurately describes the facts in the scenario. Three-year-olds, however, often reject such assertions in these scenarios (Johnson & Maratsos 1977, de Villiers 1995, Wellman et al. 2001, De Villiers & Pyers 2002, Sowalsky et al. 2009). While this might be taken as support for a conceptual change between 3 and 4 years, children's behavior might be driven by their interpretation of the linguistic stimuli, and not a paucity in their conceptual repertoire.

Lewis et al. (2017) provide support for the view that three-year-olds are (overwhelmingly) responding to the truth of the complement proposition in test sentences, rather than to the truth of the entire assertion (e.g., it would be correct to reject the assertion that Danny is at work given the scenario). Lewis and colleagues suggest that three-year-olds evaluate an utterance like (21b) with the general form *X thinks p* as a parenthetical, or an indirect assertion of *p*, which also happens to mark the evidential source of the assertion (i.e.: *p, because X thinks so*). This is compatible

with adults' interpretations, given that **assertive** verbs like *think* allow for the content of the belief ascription to take main point status in a conversation, as opposed to the ascription itself (Q: *Where's Mary?* A: *John thinks she's in her office.*) (Simons 2007, Hooper 1975, Urmson 1952). Furthermore, it may be that children are even more prone to these interpretations than adults are, since such indirect assertions are the primary input that children receive with *think* (Dudley et al. under review, see further discussion of this in §5).

The conclusions of Lewis and colleagues fit within a wider body of research suggesting that elicited-response tasks can be subject to pragmatic confounds (Siegal & Beattie 1991, Helming et al. 2016, Westra & Carruthers 2017, Hacquard & Lidz 2022, Rubio-Fernández & Geurts 2013, Białecka-Pikul et al. 2019, Baratgin et al. 2020, Salter & Breheny 2019, Hansen 2010). And the issue does not end at the borders of developmental psychology. Adults make qualitatively similar pragmatic errors of misinterpretation under certain linguistic conditions (Hansen et al. 2017, Dudley & Mascarenhas in preparation) although at quantitatively different rates. The quantitative difference in error rates that is found between adults, older children and younger children is compatible with a picture where language interpretation involves probabilistic inference, with the weights for priors shifting across development as the child gains knowledge about both the world and communicative partners (Frank & Goodman 2012, Goodman & Frank 2016, Bohn et al. 2021, Degen 2023). Such findings also suggest that linguistic stimuli could influence the responses of adults when we experimentally probe their knowledge attributions using linguistic stimuli, as a growing body of research does. Consider a study in Phillips et al. (2018). Phillips and colleagues asked adults to evaluate sentences like those in (22) after reading a narrative about Mira developing a justified true belief, and measured how long their participants took to make these evaluations.

(22) a. Mira knows she is looking at Neptune.
b. Mira thinks she is looking at Neptune.

Phillips and colleagues found that adults are quicker to evaluate *know* sentences compared to *think* sentences, and concluded that this difference in speed is due to the primacy of knowledge attributions over belief attributions, supported by the fact that the finding generalizes across a range of linguistic expressions. However, the differences in reaction times may be influenced by linguistic properties of the stimuli in (22). Related to the discussion above, sentences with **assertive** propositional attitude verbs like *think* are routinely ambiguous between a belief ascription (*X has the belief that p*) and an endorsement of the belief (*p, because X thinks so*), whereas sentences with verbs like *know* are not. Thus the slower response to *think* sentences could be due to the need to initially resolve this prevalent ambiguity before evaluating the relevant belief ascription—no such ambiguity resolution is needed to

interpret the *know* sentences. Interestingly the purported reaction time advantage for knowledge over belief is not replicated in a recent study by Bricker (2020), where the stimuli presented to participants were visual scenes and not sentences such as those in (22). Instead, language was used only to convey the task instructions. This underscores our contention that the pragmatic dimensions of linguistic stimuli and their interpretation can confound what might otherwise be taken as evidence about conceptual competence.[22]

Tasks that ask participants to evaluate linguistic stimuli offer an efficient means of recruiting their conceptual competence since they are easy to construct and easy to minimally modulate along various dimensions. But work with these stimuli should proceed keeping in mind that it directly measures linguistic understanding, that not only funnels conceptual contents in idiosyncratic ways (Glanzberg 2011, Pietroski 2010), but requires participants to recruit pragmatic processes that, as adept language users, are easy for researchers to overlook. We've reviewed in some detail compelling and canonical developmental work that does exactly this, despite diligent laudable efforts to control for confounds. But in using linguistic stimuli, measured effects may be a consequence of the influence that linguistic cognition exerts in processing expressions and understanding their intended use, and thus may not be a direct reflection of extra-linguistic cognition alone. Moving forward, we will be better able to make inferences about underlying conceptual competence if we pay attention to how participants represent the linguistic expressions that we use as stimuli.

## 5 Studying the linguistic expression

In this section, we will discuss research on understanding of the linguistic expression *know*. For reasons discussed above (in §2), we will approach this discussion through an internalist lens, treating the object of study as a mental object which influences behavior both in generating judgments about strings and/or utterances containing *know* and in producing *know*-utterances. We can study this mental object using experimental methods (including psycholinguistic and neurolinguistic approaches in addition to armchair investigations) where variables are manipulated. We can also use observational or pseudo-experimental methods such as corpus analyses (Chartrand 2022) to support hypotheses about the properties of expressions like *know*. But as we'll see, such analyses prove more useful for the study of the linguistic expression *know* than the concept KNOW, given practical considerations pertinent to the range of different behaviors that each one generates. We'll begin by considering some relevant properties of *know* that the descriptive and formal literatures in linguistics

22 Our thanks to an anonymous reviewer for drawing our attention to this distinction.

have uncovered, and then review some experimental methods and corpus analyses that inform both adults' understanding of the linguistic expression and how children's understanding of it develops. Throughout, it will prove useful to use *think* as a foil to *know*.

### 5.1 Properties of *know* and *think*

We start with two central properties of *know* that differentiate it from the closely related linguistic expression *think*: **veridicality** and **factivity**. In the formal linguistics literature, veridical verbs are those that **entail** the truth of the proposition expressed by their complement clause. Factive verbs are those that **presuppose** the truth of the proposition expressed by their complement clause.

For the sake of interdisciplinary discussions, we advocate for maintaining the distinction between veridicality and factivity in the formal linguists' sense. Veridicality regards the behavior of *know* in basic declarative clauses, while factivity pertains to **entailment-cancelling contexts**—for example, when the complement falls under the scope of negation, modals, conditionals or questions. Factivity was originally identified in linguistics to identify the presuppositional behavior of verbs like *know*. The broader use of the term "factive" outside of linguistics describes the veridicality of *know*-ascriptions.[23] We advocate returning to the term veridical, whose etymological history is longer and its intuitive meaning more consistent with the terminological use.

Beyond this basic distinction between factive and non-factive verbs, the formal linguistics literature also distinguishes between different kinds of factive verbs. Karttunen (1971) points out that the presupposition triggered by verbs like *discover* are more easily cancellable in certain contexts than the presuppositions of verbs like *regret*. Verbs in the former group are often called cognitive factives or semi-factives while verbs in the latter group may be called true factives or emotive factives. There is also some disagreement about how to capture these two kinds of factives, whether they should be treated as different (Simons 2007, Abrusán 2011) or the same (Abbott 2006). And, more generally, there is much debate over whether presuppositions should be directly encoded as part of the lexical meaning of a presupposition trigger or whether they might arise downstream through the interaction of a presupposition trigger's meaning and general pragmatic principles (Schlenker 2021). In recent years, experimental work has suggested that adults' judgments of factive presupposition are less categorical than once thought (White & Rawlins 2018, de Marneffe et al. 2019, Degen & Tonhauser 2022).

23 See Nagel 2017 and Egré 2008 for more discussion of how linguists and philosophers each use the term "factive".

*Know* is also distinguished from *think* by the kinds of syntactic environments that it occurs in. *Know* and other factive verbs are **responsive** in that they take both declarative and interrogative clauses (Lahiri 2002). In contrast, non-factive verbs tend to be non-responsive in that they either embed only interrogatives (e.g., *wonder*) or only declaratives (e.g., *think*). The link between the syntactic properties of verbs like *know* and their semantic properties was observed long ago (Hintikka 1975, Ginzburg 1995) but it remains to be seen how the link should be formally captured (Anand & Hacquard 2014, Roberts 2021, White 2021, Spector & Egré 2015, Uegaki 2015, Ciardelli et al. 2015).

### 5.2 Language acquisition

Returning again to the question of development, studies in language acquisition have pursued when and how children master the subtleties of *know* and learn to distinguish it from *think*.

#### 5.2.1 Experimental methods

One set of researchers have focused on verbs like *know*, asking when children understand these verbs' subtle semantic properties, as compared to other verbs. This literature suggests that children do not understand the difference between *know* and *think* before four years (Macnamara et al. 1976, Abbeduto & Rosenberg 1985, Moore & Davidge 1989, Moore et al. 1989) or even well into the school years (Harris 1975, Hopmann & Maratsos 1978, Scoville & Gordon 1980, Falmagne et al. 1994, Schulz 2003, Léger 2008). Some comprehension studies ask when children understand that *know* statements are stronger, or more informative, than *think* statements (perhaps because *know* statements express greater certainty from the speaker). As one example, Moore and colleagues gave children a search task where they must find a hidden toy on the basis of two (conflicting) cues (Moore & Davidge 1989, Moore et al. 1989). One puppet would utter *I know it's in the red box* while another would say *I think it's in the blue box*, and children were expected to pick the red box in order to demonstrate that they understand the greater strength of the *know* claim. Moore and colleagues found that three-year-olds were unable to reliably search in the location described by the puppet who *knows* as opposed to the puppet who *thinks* or *guesses*, while four-year-olds could. However, while these results demonstrate a developmental change between 3 and 4, it is not clear that the change is rooted in understanding of *know* vs. *think*, as opposed to either their understanding of general pragmatic principles involved in comparing contrasting claims, or the way they use information gathered from testimony to guide their own

behavior. Again, we see that even for those directly studying language, assumptions about pragmatics and the use of language confound the informativeness of results.

More directly relevant to factivity, some comprehension studies ask whether children understand that verbs like *know* allow the truth of their complements to project out of entailment-cancelling contexts, such as negation. These studies ask children to infer whether *p* is true when they hear utterances like *x doesn't know that p* and *x doesn't think that p*. Many of these studies suggest that younger children do not reliably infer that the complement to *know* is true, thus failing to differentiate it from *think* (Harris 1975, Hopmann & Maratsos 1978, Scoville & Gordon 1980, Léger 2008). However, the methods sometimes required literacy (e.g., filling out answers in a test booklet) or involve convoluted scenarios (e.g., a game show with a fortune teller), which might preclude younger children from reliably demonstrating their understanding of the verbs.

A different literature focuses on errors that children make with *think* until about 4 years of age, which is when they start to pass elicited-response false belief tasks with linguistic stimuli (described above in §4, see the example in (21)). Recall, these studies suggested that three-year-olds incorrectly reject a true *think* sentence that ascribes a false belief (Johnson & Maratsos 1977, de Villiers 1995, Wellman et al. 2001, De Villiers & Pyers 2002, Sowalsky et al. 2009). Interestingly, this would be the expected response if three-year-olds assumed that *think* was a veridical (or even factive) verb.

Notice the conflict between these two bodies of results. The first would lead us to conclude that three-year-olds have failed to learn the factivity of *know*, because they treat it the same way that adults treat *think*. The second would lead us to conclude that three-year-olds have mistakenly learned that *think* is veridical, because they treat it the same way that adults might treat *know*. Is it really possible that children's early failures are so absolute that they have learned the verbs backwards (and if so, how would they ever overcome this)? Or can we find some alternative explanation of the pattern of results?

Dudley and colleagues have tried to re-assess three-year-olds' understanding of verbs like *think* and *know* to settle the conflict raised by the literature (Dudley et al. 2015, in preparation). Their working assumption is that previous methods were insufficient for the question at hand because those methods (i) involved dependent measures that are ill-suited for children as young as 3 years, or (ii) did not examine entailment-cancelling contexts—linguistic environments that distinguish the two verbs on the basis of their (non-)factivity. As a result, their series of studies builds upon methods focusing on entailment-cancelling contexts (Scoville & Gordon 1980) and updates them to a more age-appropriate dependent measure for preschoolers (Moore & Davidge 1989). Participants are asked to find a hidden toy in one of two boxes after hearing a propositional attitude report, as in (23–26).

(23) Lambchop **doesn't** know that it's in the blue box. (Target response: search in the blue box)

(24) Lambchop **doesn't** think that it's in the blue box. (Target response: search in the other box)

(25) Lambchop knows that it's in the blue box. (Target response: search in the blue box)

(26) Lambchop thinks that it's in the blue box. (Target response: search in the blue box)

Children's searching behavior allows for inferences about their understanding of the verbs' (non-)factivity. If they understand a verb to be factive, then they should search for the toy in the box labelled in the complement clause, regardless of the presence or absence of an entailment-cancelling operator such as negation.[24] The results suggested that the non-factivity of *think* is acquired relatively early.[25] With respect to *know*, there was individual variation in children's understanding of its factivity where some children had mastered it by their 3rd birthday while others had not, even at 4.5 years. Dudley and colleagues took this to suggest that (i) the non-factivity of *think* is easy to discern, while (ii) the factivity of *know* is more challenging, but (iii) it is ultimately discovered between 2–5 years of age. Notice that this method does not fall victim to the pragmatic confounds raised for studies above. For example, both the belief report interpretation and the endorsement interpretation of (26) would lead children to the same searching behavior (because they are given no reason to mistrust Lambchop, or the speaker who is giving them useful clues to find the toys they desire).

Given these findings, the noted tension in the acquisition of *think* versus *know* can be explained (at least in part) by a failure to attend to the pragmatics of language use. That a *know*-ascription takes a true complement while a *think*-ascription need not is a property of those expressions. But that a speaker can intend to use the information imparted by this difference to indicate the reliability of a source of information is a pragmatic feature of language use. Children's ability to recognize that a speaker

24 By definition, this kind of task requires the participant to accommodate the presupposition triggered by *know*. With a well-designed study, this has the desirable consequence that the bar for success is quite high: participants succeed if and only if they both understand what is presupposed and are able to accommodate it online. There may be many reasons why children (and even adults) might fail to accommodate.

25 Of course a method based on presupposition projection can only rule out the possibility that children have a factive representation for *think*. However it is still an open question whether children ever entertain a non-factive but veridical representation. This would be inconsistent with the generalization that all doxastic factives are also veridicals (Anand & Hacquard 2014) and would require demonstrating that there is evidence available to the learner allowing them to un-learn the veridicality of *think*.

can use a single expression with various intentions, and their ability to identify what those intentions are, is a capacity that develops alongside the acquisition of lexical items. Failing to account for this developing pragmatic competence invites misleading analyses.

### 5.2.2 Corpus methods

Another body of work looks to corpora to investigate children's acquisition of verbs like *know* and *think*. Typically, these investigations examine either children's use of the verbs to conclude something about their understanding of the verbs, or they examine the ways that adults use the verbs around children to understand what evidence is available to children as they learn about the verbs' meanings. We'll first review research on children's own productions and then turn to input studies.

Going back several decades, corpus studies have demonstrated that *know* and *think* are produced early in development, but this data has also been used to argue that the earliest uses are not reflective of adult-like competence with the verbs. Before we review more recent research which challenges this perspective, we will discuss some of the older findings.

One of the first studies to examine children's productions of attitude verbs both longitudinally (across time) and cross-sectionally (across individuals) was conducted by Shatz et al. (1983) using the Abe corpus (one target child, followed from 28 months to 5 years, Kuczaj 1978, MacWhinney 2000) and a cross-sectional corpus of 30 two-year-olds (Hoff-Ginsberg 1981). Shatz and colleagues found that *know* was the doxastic verb used most frequently by children, with the first uses emerging around 2.5 years. In contrast, *think* was found to occur much less often and to emerge a few months later than *know*. According to Shatz and colleagues, *think* is used most often to describe others' beliefs or to modulate the force of an assertion (e.g., hedging), while *know* is used to draw an interlocutor's attention or to request information. Despite this nuanced picture, Shatz and colleagues conclude that these early productions are not indicative of a flexible adult-like understanding of the verbs because of the prevalence of particular forms such as *I don't know*, which may be understood idiomatically.[26] However, Shatz and colleagues (along with other authors at the time, such as Bartsch & Wellman 1995) placed a high threshold for the evidence required to conclude children have a mature understanding of the verbs. They sought utterances which unambiguously contrasted verbs like *know* and *think* as in (27), even though such reports of false beliefs are rare in adult language.

26 See also Diessel & Tomasello 2001 for similar arguments about early uses of *think* and Bretherton & Beeghly 1982 for parental report data which corresponds to the corpus data.

(27) Before I thought this was a crocodile; now I know it's an alligator. (Shatz et al. 1983: pg. 309)

However, a new look at young children's productions suggests that older studies underestimated children's understanding of the verbs and their sophistication in using the verbs. New analysis by Harris et al. (2017b) suggests that *I don't know* utterances need not be idiomatic and can be reflective of a deeper understanding of the verb's meaning. Children use utterances like these to respond to information-seeking questions before they are 3 years old, which is fully compatible with adult-like usage. Furthermore, it would be reasonable if there was a difference in how often children vs. adults use such a sentence to express their ignorance, given that children are less likely to possess information than their elders. This kind of new analysis suggests that the gulf between children's and adults' competence is not necessarily so wide, reflecting a quantitative difference in the kinds of messages they wish to convey, not a qualitative difference in their understanding of the relevant verb meanings.

We turn now to corpus analyses of adult speech to children (child-ambient speech). It can be informative to study adult productions of *know* for non-developmental reasons (see discussion above in §3), but why is it useful to the study of language development? Answering this question requires some background on the nature of language acquisition. Every approach to language acquisition recognizes the importance of **input** to the process (Lidz & Gagliardi 2015). At the end of the day, you need some information to tell you that a concept like APPLE is recruited by the word *apple* in one language, *pomme* in another, and *alma* in yet another. We still know relatively little about how the language learner takes in such information to generate adult linguistic competence. But if we conceive of the process as a function which goes from certain inputs to certain outputs, we can observe the inputs and outputs and try to infer what kind of function would generate them. On a general level, this research program has illustrated that both **quantity** and **quality** of input matter for language development. Famously, Hart & Risley (1995) found that children from different socioeconomic backgrounds heard different amounts of words in the preschool years and that these differences were correlated with children's performance on language measures when they entered school. But even holding overall frequency constant, more qualitative aspects of input such as complexity of syntactic constructions and frequency of joint attention can also influence aspects of language development such as the size of children's vocabulary (Hoff 2006), the diversity of syntactic constructions in their speech (Huttenlocher et al. 2010) and even the age at which children learn a particular word (Hoff & Naigles 2002). So what can investigations of input (and their relationship to output) tell us about *know*?

Attitude verbs like *know* and *think* are relatively frequent in naturalistic speech, considering the enormous number of other words that could be used. Attitude

verbs are produced in around 10% of utterances cross-linguistically and across conversational contexts (Adrián et al. 2007, Furrow et al. 1992, Slaughter et al. 2007, Tardif & Wellman 2000). *Know* and *think* are two of the most frequent, occurring approximately 53 thousand and 35 thousand times respectively in the CHILDES database (MacWhinney 2000), which equates to about 1–2 times per 100 utterances. The frequency of *know* has been discussed to illustrate the ubiquity of it and therefore the potential simplicity of it (Nagel 2017).

But such a claim about the simplicity of *know* belies the difficulty of the acquisition task for children. To see this, compare *know* to *apple*. Over 90% of children understand *apple* by 18 months of age and produce it by 23 months, according to normed parental reports of receptive and productive vocabulary (Dale & Fenson 1996, Frank et al. 2017). In contrast, *know* is only understood by a majority of children about 2 years later. However, if we consider the frequency of the words, *know* occurs 40 times more often (even in the context of parent-child conversations) than *apple* does (MacWhinney 2000). If frequency were all that mattered, why wouldn't children master *know* before *apple*? One possibility is that *know* is more challenging because of the nature of the concept that it must be linked to (see discussion above in §4). But it is also hard to observe the referent of a propositional attitude report in the wild because mental states seem to lack reliable physical correlates, such as those exemplified in many a Rodin sculpture (Gleitman 1990). As a result, Lila Gleitman (and many inspired by her) have taken the linguistic input with these verbs to be particularly important to unlocking the acquisition puzzle (see for example Gillette et al. 1999, Gleitman et al. 2005, Papafragou et al. 2007a). We can see this even as adults, when we are given a unknown verb in a known linguistic context:

(28) Alex daxes that Lewis is supposed to arrive later in the afternoon.

(29) But Alex doesn't dax what the exact time will be.

(30) Do you dax what time Lewis will arrive?

Upon hearing (28), you may have the intuition that *dax* means something like *think*, *know* or *say*. After hearing (29), you would be able to rule out *think*, likely due to your implicit understanding of the syntactic complements that the verbs take. Furthermore, upon hearing (30), you would be able to further eliminate *say* from consideration, likely due to your understanding of the potential discourse moves we can use the verbs to make.

In a corpus study, Dudley and colleagues investigated whether there are such cues in child-ambient speech that would give away what *know* means when compared to *think* (Dudley et al. 2017, under review). They asked whether children only hear *know* and *think* in utterances like (28) that are syntactically indistinguishable from each other (and, if so, what other information would they use to learn about the semantics of the verbs?). Or do children also hear utterances like (29–30)

that reveal the differences between *know* and *think*? Dudley and colleagues used the Gleason corpus, a cross-sectional corpus of input to 30 children between 2–5 years participating in play and mealtime interactions with their parents (Masur & Gleason 1980, MacWhinney 2000). *Know* and *think* each occurred in about 3% of the child-ambient utterances in this corpus. Dudley and colleagues examined the role of two kinds of linguistic cues in differentiating *know* from *think*: (i) direct cues to the factivity/veridicality contrast via the discourse status of the complement and (ii) indirect cues via discourse function. Cues of type (i) involve evidence of the entailment patterns and presuppositional behavior of the two verbs (or lack thereof). While this type of cue would be directly useful to mapping the different meanings of the two verbs, the researchers did not find that it reliably distinguished the two verbs in the input. While *think* can in principle occur with complements that express false propositions (e.g., *He thinks that Berlin is in France*), it was mostly used to perform indirect assertions (e.g., *I think it's time for bed*) in children's experience. Similarly, while *know* presupposes the truth of its complement, it was often used "informatively" to provide new information to a listener, which would not distinguish it from *think* (e.g., *Did you know that your grandmother was a seamstress?*, see also Spenader 2003 for similar results in adult-to-adult speech). Relevant to cues of type (ii), the researchers found that the most prevalent uses of *think* were indirect assertions whereas *know* was mostly used to ask or answer questions (e.g., *Do you know where my keys are? I don't know.*). Overall, these results suggest that there is information in the linguistic environments of the two verbs which could be used to learn the difference between the verbs. In particular, cues from the discourse functions that the two verbs are used to achieve (and the particular syntactic frames used to realize these discourse functions) are useful under pragmatic-syntactic bootstrapping, a learning strategy which is also useful for acquiring other semantic distinctions in the attitude verb domain (such as *think* vs. *want*, see Hacquard & Lidz 2022 for further discussion). But what should we make of these suggestive results from the corpus study? It is only an observational study, or pseudo-experimental at best given that it compares how *know* surfaces in the corpus compared to *think*. We need further methods to test whether these differences between *know* and *think* are related with children's understanding of the two verbs.

Correlational methods are a useful next step because they establish that there is a relation between the input and ultimate understanding, even if they cannot establish a causal link.[27] Within the domain of attitude verbs, we know that children who

27 A relation between the two is also consistent with a causal role in the opposite direction to the one we seek. For example, children may get more of the hypothesized critical input because their understanding was already more mature than their peers and thus they lead adults into more complex conversations than their peers. Ultimately, experimental methods like training studies are needed to test causal hypotheses.

hear more attitude verbs in the input also produce more attitude verbs themselves (Booth et al. 1997, Jenkins et al. 2003). Additionally, a correlational study in Howard et al. (2008) analyzing coarse-grained aspects of the input found that some aspects are related to children's performance on the Moore & Davidge (1989) task. First, hearing *know* in more questions was related to better performance, while hearing *know* in more declaratives was related to worse performance. Second, hearing *know* occur with more first person subjects was related to worse performance.

Building on these findings, Dudley and colleagues have asked how fine-grained aspects of linguistic input are correlated with children's understanding of the factivity of *know* as measured via understanding of presupposition projection. To accomplish this, they integrated the corpus methods (Dudley et al. 2017, under review) and behavioral tasks (Dudley et al. 2015) discussed above, with initial results reported in (Dudley 2017). They found that both quantity and quality of the input were related to better understanding that *Chris doesn't know that the toy is in the blue box* presupposes that the toy is indeed in the blue box. On the role of quantity, they found that children who heard the verbs more were better at distinguishing the verbs in the behavioral task. On the role of quality, they found two interesting relations. On one side, children whose parents provided better discourse function cues (by using *think* in indirect assertions and by using *know* to respond to information-seeking questions) better understood the distinction between the verbs. On the other side, children whose parents provided more cues about the discourse status (by using *think* to express false beliefs more than *know* or by using *know* only when the complement proposition was presupposed) were more likely to treat the verbs indistinguishably.

A note of caution about corpus analyses should be offered, based on this line of research. Corpora are bodies of observational data, and as such they do not generate a perfect record of the range of possible behaviors. There are many reasons (intuitive but also less so) why people in everyday conversation do not say everything that they could say: it may not be relevant nor appropriate nor useful in a particular context. As a result, the inferences that corpus work licenses are different from those licensed by experimental work. Care should be taken in limiting the extent/range of conclusions that we draw from corpus methods, and efforts should be made to use corpora that are as representative as possible. Take the work of Dudley and colleagues: although they found the same general pattern of results across their own corpus and pre-existing corpora, the specifics within each corpus were different to each other. Overall, both bodies of data have allowed us to form the hypothesis that children learn the veridicality (and thus infer the factivity) of *know* by hearing the verb used to achieve interrogative discourse functions, such as asking questions and answering them with expressions of ignorance. However, if this question had just been probed in the Gleason corpus, where the primary use of *know* was in information-seeking questions such as *Do you know what time it is?*, they would

have formed a hypothesis about the specific role of indirect requests for information. In contrast, if this question had just been probed in their own corpus, where the primary use of *know* was to express ignorance by saying *I don't know* in response to questions, then they would have formed a hypothesis about the specific role of providing answers to information-seeking questions. But by examining both in conjunction, we likely achieve something closer to the truth.

To turn back to the discussion in §3-4, how do these language acquisition findings speak to issues that have been raised in those literatures? They demonstrate that language acquisition is a nuanced and extended process, and that children may not always understand what we say to them in the way we intend it. They are still in the process of setting their priors for pragmatic inference, by building their body of world knowledge and accruing a set of experiences in communicative settings with communicative partners. As a result, we should take care when we present children with linguistic stimuli in order to test some other competence. Unexpected behavior may be due to children's immature competence in the conceptual domain, to their developing linguistic representations, or even to what they assume the experimenter wants to achieve in the experimental setting.

In addition, we should be careful not to draw quick conclusions. By reviewing only some of the studies in this literature, it would seem that *know* is mastered early or plays a more prominent role in children's mental state language than *think* (Nagel 2013). However, this discounts that there may be practical reasons why children utter *know* sentences more often than *think* sentences (Harris et al. 2017b); and also that there are linguistic properties of *know*—such as its presuppositions—that present an acquisition puzzle which does not arise for *think* (Dudley et al. 2015, 2017).

## 6 Conclusion

The nature of knowledge is tied to the human ability to represent knowledgeable minds, both in terms of the concepts like KNOW recruited by our conceptual competence, and expressions like *know* that somehow recruit these representations in the communicative use of a language. As a result, the use of language in methods that investigate knowledge and the mind-reading systems that represent knowledgeable minds, and indeed epistemic language itself, prove invaluable. We've reviewed how this use manifests in various methods: the method of cases, elicited-response tasks, corpus methods, and others. We've highlighted throughout that linking assumptions are required to yield verdicts about the import of participant behavior to these various targets of inquiry. We argued that a pair of optimistic assumptions that would provide the needed links are undermined by pervasive phenomena in natural language.

The first assumption—that the meaning of a natural language expression determines the concept/intension it expresses (even on a particular use)—is undermined

by the ubiquity of polysemy, which illustrates that a single meaning maps to many concepts/intensions. The second assumption—that a speaker's utterances (or a listener's understanding of another's utterances) are solely reflective of either their underlying linguistic or extra-linguistic cognition—is in tension with the ubiquitous fact that speakers use expressions for a wide-variety of purposes in communication. This tells against accepting participants' behavior as directly reflective of what they are thinking, even in the rarefied experimental context. Of course, good scientific hygiene requires accounting for confounds in the methods utilized, and the research highlighted in §§4–5 surely respects this mandate in inventive and laudable ways. But the point we've emphasized is that, precisely because human linguistic communication often seems effortless, special care needs to be placed in attending to the various, often subtle, pragmatic aspects of language use. This point is particularly salient in the developmental literature because we still do not fully understand the developmental course of pragmatic abilities. Since this capacity is developing in parallel with both children's syntactic-semantic competence and their expanding extra-linguistic conceptual repertoire, the ontogenetic trajectory of any of these systems can be responsible for the non-adult-like behavior of children. While researchers are often sensitive to the role of children's developing syntactic-semantic competence, far less attention is paid to their pragmatic understanding. We show how this worry arises in corpus analysis as well. While corpora can indicate what inputs are available to children in their language learning environment, the informativeness of this input is, for the child, mediated by their pragmatic competence. A linguistic stimulus may have a different perlocutionary effect on a child—even if a developing child has the full conceptual repertoire of an adult—simply because they have less experience in how adults use language to convey communicative intentions.

These points are of particular importance to an ordinary language philosopher that takes considerations of language use as essential to methodology. That methodological mandate compels us to consider how individuals process the utterances they encounter in deciphering the communicative intentions of speakers. As such, assumptions about linguistic meaning and use constrain what those processes are like, which can mislead as much as they can inform. Failing to attend to the import of polysemy and pragmatics in these methods has the potential to frustrate that insightful mandate, in the varied and subtle ways we've highlighted here. Language offers us a window into the mind, giving us access to the conceptual contents that natural language expressions somehow recruit. But windows can contain broken and opaque panes, obscuring the objects they frame. Understanding, and attending to the ways that language use and comprehension can distort the underlying thoughts our utterances express is required if our language-centric experimental methods are to elucidate how we represent the world.

## References

Abbeduto, Leonard & Sheldon Rosenberg. 1985. Children's knowledge of the presuppositions of know and other cognitive verbs. *Journal of Child Language* 12(03). 621–641.

Abbott, Barbara. 2006. Where have some of the presuppositions gone. In Gregory Ward & Betty J Birner (eds.), *Drawing the Boundaries of Meaning: Neo-Gricean Studies in Pragmatics and Semantics in Honor of Laurence R. Horn*, 1–20. John Benjamins Publishing Company.

Abrusán, Márta. 2011. Predicting the presuppositions of soft triggers. *Linguistics and Philosophy* 34(6). 491–535.

Adrián, Juan E, Rosa Ana Clemente & Lidón Villanueva. 2007. Mothers' use of cognitive state verbs in picture-book reading and the development of children's understanding of mind: a longitudinal study. *Child Development* 78(4). 1052–1067.

Anand, Pranav & Valentine Hacquard. 2014. Factivity, Belief and Discourse. In Luka Crnič & Uli Sauerland (eds.), *The Art and Craft of Semantics: A Festschrift for Irene Heim*, vol. 1, 69–90. Cambride, MA: MIT Working Papers in Linguistics.

Annis, David B. 1978. A contextualist theory of epistemic justification. *American Philosophical Quarterly* 15(3). 213–219.

Apresjan, J. 1974. Regular polysemy. *Linguistics* 14. 5–32.

Austin, John Langshaw. 1975. *How to Do Things with Words*. Oxford University Press.

Baillargeon, Renée, David Buttelmann & Victoria Southgate. 2018. Invited commentary: Interpreting failed replications of early false-belief findings: Methodological and theoretical considerations. *Cognitive Development* 46. 112–124.

Baratgin, Jean, Marion Dubois-Sage, Baptiste Jacquet, Jean-Louis Stilgenbauer & Frank Jamet. 2020. Pragmatics in the false-belief task: Let the robot ask the question! *Frontiers in Psychology* 11. 3234.

Baron, Robert S, Joseph A Vandello & Bethany Brunsman. 1996. The forgotten variable in conformity research: Impact of task importance on social influence. *Journal of Personality and Social Psychology* 71(5). 915.

Barrett, H Clark, Tanya Broesch, Rose M Scott, Zijing He, Renée Baillargeon, Di Wu, Matthias Bolz, Joseph Henrich, Peipei Setoh, Jianxin Wang et al. 2013. Early false-belief understanding in traditional non-western societies. *Proceedings of the Royal Society B: Biological Sciences* 280(1755). 20122654.

Bartsch, Karen & Henry M. Wellman. 1995. *Children talk about the mind*. Oxford University Press.

Baz, Avner. 2017. *The Crisis of Method in Contemporary Analytic Philosophy*. Oxford University Press.

Begus, Katarina & Victoria Southgate. 2012. Infant pointing serves an interrogative function. *Developmental Science* 15(5). 611–617.

Berwick, Robert C, Paul Pietroski, Beracah Yankama & Noam Chomsky. 2011. Poverty of the stimulus revisited. *Cognitive Science* 35(7). 1207–1242.

Białecka-Pikul, Marta, Magdalena Kosno, Arkadiusz Białek & Marta Szpak. 2019. Let's do it together! the role of interaction in false belief understanding. *Journal of Experimental Child Psychology* 177. 141–151.

Birch, Susan AJ & Paul Bloom. 2003. Children are cursed: An asymmetric bias in mental-state attribution. *Psychological Science* 14(3). 283–286.

Birch, Susan AJ & Paul Bloom. 2007. The curse of knowledge in reasoning about false beliefs. *Psychological Science* 18(5). 382–386.

Bohn, Manuel, Michael Henry Tessler, Megan Merrick & Michael C Frank. 2021. How young children integrate information sources to infer the meaning of words. *Nature Human Behaviour* 5(8). 1046–1054.

Bolinger, Renee. 2020. Contested slurs: Delimiting the linguistic community. *Grazer Philosophische Studien* 97(1). 11–30.

Booth, James R, William S Hall, Gregory C Robison & Su Yeong Kim. 1997. Acquisition of the mental state verb know by 2-to 5-year-old children. *Journal of Psycholinguistic Research* 26(6). 581–603.

Bourget, David & David J Chalmers. 2014. What do philosophers believe? *Philosophical Studies* 170. 465–500.

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

Bricker, Adam Michael. 2020. The neural and cognitive mechanisms of knowledge attribution: An eeg study. *Cognition* 203. 104412.

Bromberger, Sylvain. 2011. What are words? comments on kaplan (1990), on hawthorne and lepore, and on the issue. *The Journal of Philosophy* 108(9). 486–503.

Burge, Tyler. 2003. Psychology and the Environment: Reply to Chomsky. In Martin Hahn & B. Ramberg (eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, 451–471. MIT Press.

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

Call, Josep & Malinda Carpenter. 2001. Do apes and children know what they have seen? *Animal Cognition* 3(4). 207–220.

Cappelen, Herman. 2020. Conceptual Engineering: The Master Argument. In Herman Cappelen, David Plunkett & Alexis Burgess (eds.), *Conceptual Engineering and Conceptual Ethics*, Oxford: Oxford University Press.

Cappelen, Herman & Matthew McKeever. 2022. In defense of ordinary language philosophy. *Metaphilosophy* 53(2-3). 221–237.

Carlson, Stephanie M & Louis J Moses. 2001. Individual differences in inhibitory control and children's theory of mind. *Child Development* 72(4). 1032–1053.

Chartrand, Louis. 2022. Modeling and corpus methods in experimental philosophy. *Philosophy Compass* e12837.

Chomsky, Noam. 1965. *Aspects of a Theory of Syntax*. MIT Press.

Chomsky, Noam. 1980. On cognitive structures and their development: a reply to piaget. In Massimo Piattelli-Parini (ed.), *Language and Learning: The Debate Between Jean Piaget and Noam Chomsky*, Harvard University Press.

Chomsky, Noam. 1986. *Knowledge of Language*. Praeger Publishing.

Chomsky, Noam. 2000. *New horizons in the study of language and mind*. Cambridge University Press.

Ciardelli, Ivano, Jeroen Groenendijk & Floris Roelofsen. 2015. On the semantics and logic of declaratives and interrogatives. *Synthese* 192(6). 1689–1728.

Clark, Andy & David Chalmers. 1998. The Extended Mind. *Analysis* 58(1). 7–19.

Clark, Brady. 2020. The evolvability of words: On the nature of lexical items in minimalism. *Frontiers in Psychology* 10.

Clements, Wendy A & Josef Perner. 1994. Implicit understanding of belief. *Cognitive Development* 9(4). 377–395.

Cohen, Stewart. 1988. How to be a fallibilist. *Philosophical Perspectives* 2. 91–123.

Cohen, Stewart. 1999. Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives* 13. 57–89.

Crain, Stephen & Paul Pietroski. 2001. Nature, nurture, and universal grammar. *Linguistics and Philosophy* 24. 139–186.

Culicover, Peter W & Ray Jackendoff. 2010. Quantitative methods alone are not enough: Response to gibson and fedorenko. *Trends in Cognitive Sciences* 14(6). 234–235.

Dale, Philip S & Larry Fenson. 1996. Lexical development norms for young children. *Behavior Research Methods, Instruments, & Computers* 28(1). 125–127.

Davidson, Donald. 1967. The Logical Form of Action Sentences. In Nicholas Rescher (ed.), *The logic of decision and action*, University of Pittsburgh Press.

Davidson, Donald. 1973b. Radical interpretation. *Dialectica* 27. 313–328.

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

Degen, Judith. 2023. The rational speech act framework. *Annual Review of Linguistics* 9. 519–540.

Degen, Judith & Judith Tonhauser. 2022. Are there factive predicates? an empirical investigation. *Language* 98(3). 552–591.

DeRose, Keith. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research* 52(4). 913–929.

DeRose, Keith. 1995. Solving the skeptical problem. *Philosophical Review* 104(1). 1–52.

Devitt, Michael. 2006. *Ignorance of language*. Oxford University Press.

Diessel, Holger & Michael Tomasello. 2001. The acquisition of finite complement clauses in english: A usage based approach to the development of grammatical constructions. *Cognitive Linguistics* 12. 97–141.

Dinges, Alexander & Julia Zakkou. 2021. Much at stake in knowledge. *Mind & Language* 36(5). 729–749.

Dörrenberg, Sebastian, Hannes Rakoczy & Ulf Liszkowski. 2018. How (not) to measure infant theory of mind: Testing the replicability and validity of four non-verbal measures. *Cognitive Development* 46. 12–30.

Dudley, Rachel. 2017. *The role of input in discovering presupposition triggers: Figuring out what everybody already knew*: University of Maryland dissertation.

Dudley, Rachel. 2018. Young children's conceptions of knowledge. *Philosophy Compass* 13(6). e12494.

Dudley, Rachel, Valentine Hacquard & Jeffrey Lidz. in preparation. Preschooler's understanding of the projection behavior of know: A replication.

Dudley, Rachel & Mascarenhas. in preparation. Pragmatic effects in adults' evaluation of false belief reports.

Dudley, Rachel, Naho Orita, Valentine Hacquard & Jeffrey Lidz. 2015. Three-year-olds' understanding of know and think. In *Experimental Perspectives on Presuppositions*, 241–262. Springer.

Dudley, Rachel, Meredith Rowe, Valentine Hacquard & Jeffrey Lidz. 2017. Discovering the factivity of "know". In *Semantics and Linguistic Theory*, vol. 27, 600–619.

Dudley, Rachel, Meredith Rowe, Valentine Hacquard & Jeffrey Lidz. under review. The difficulty of knowing: acquiring the contrast between know and think.

Dudley, Rachel & Christopher A. Vogel. 2024. Knowledge, the concept KNOW and the word *know*: considerations from polysemy and pragmatics. *Synthese* 203(17).

Dummett, M. 1978. What do I Know When I Know a Language. *Universitas Regia Stockholmensis* Reprinted in Dummett, M. (1993) *Seas of Language* (pp. 94–105).

Dummett, Michael. 1986. A nice derangement of epitaphs: Some comments on Davidson and Hacking. In E. Lepore (ed.), *Truth and interpretation: Perspectives on the philosophy of donald davidson*, 459–476. Basil Blackwell Oxford.

Dummett, Michael. 1993. *The seas of language*. Oxford University Press.

Egré, Paul. 2008. Question-embedding and factivity. *Grazer Philosophische Studien* 77(1). 85–125.

Falkum, Ingrid Lossius & Agustin Vicente. 2015. Polysemy: Current Perspectives and Approaches. *Lingua* 100(157). 1–16.

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

Francis, Kathryn, Philip Beaman & Nat Hansen. 2019. Stakes, scales, and skepticism. *Ergo: An Open Access Journal of Philosophy* 6.

Frank, Michael C, Mika Braginsky, Daniel Yurovsky & Virginia A Marchman. 2017. Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language* 44(3). 677–694.

Frank, Michael C & Noah D Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998–998.

Furrow, David, Chris Moore, Jane Davidge & Lorraine Chiasson. 1992. Mental terms in mothers' and children's speech: Similarities and relationships. *Journal of Child Language* 19(03). 617–631.

Gertler, Brie. 2012. Understanding the internalism-externalism debate: what is the boundary of the thinker? *Philosophical Perspectives* 26. 51–75.

Gettier, Edmund L. 1963. Is Justified True Belief Knowledge? *Analysis* 23(6). 121–123.

Gibson, Edward & Evelina Fedorenko. 2013. The need for quantitative methods in syntax and semantics research. *Language and Cognitive Processes* 28(1-2). 88–124.

Gibson, Edward A & Evelina G Fedorenko. 2010. Weak quantitative standards in linguistics research. *Trends in Cognitive Science* 14(6). 233–234.

Gillette, Jane, Henry Gleitman, Lila Gleitman & Anne Lederer. 1999. Human simulations of vocabulary learning. *Cognition* 73(2). 135–176.

Ginzburg, Jonathan. 1995. Resolving questions, I & II. *Linguistics and Philosophy* 18(6). 459–527, 567–609.

Glanzberg, Michael. 2011. Meaning, Concepts, and the Lexicon. *Croatian Journal of Philosophy* 11(1). 1–29.

Glanzberg, Michael. 2014. Explanation and partiality in semantic theory. *Metasemantics: New Essays on the Foundations of Meaning* 259–292.

Glanzberg, Michael. 2018. Lexical meaning, concepts, and the metasemantics of predicates. In D. Ball & B. Rabern (eds.), *The science of meaning: Essays on the metatheory of natural language semantics*, 197–225. Oxford University Press.

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

Gleitman, Lila R, Kimberly Cassidy, Rebecca Nappa, Anna Papafragou & John C Trueswell. 2005. Hard words. *Language Learning and Development* 1(1). 23–64.

Goldman, Alvin. 2005. Kornblith's naturalistic epistemology. *Philosophy and Phenomenological Research* 71(2). 403–410.

Goodman, Noah D & Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences* 20(11). 818–829.

Gopnik, Alison & Peter Graf. 1988. Knowing how you know: Young children's ability to identify and remember the sources of their beliefs. *Child Development* 1366–1371.

Gotham, Matthew. 2017. Composing criteria of individuation in copredication. *Journal of Semantics* 34(2). 333–371.

Gotham, Matthew. 2022. Property inheritance, deferred reference and copredication. *Journal of Semantics* 39. 87–116.

Grant, Janie Busby & Thomas Suddendorf. 2011. Production of temporal terms by 3-, 4-, and 5-year-old children. *Early Childhood Research Quarterly* 26(1). 87–95.

Grindrod, Jumbly, James Andow & Nat Hansen. 2019. Third-person knowledge ascriptions: A crucial experiment for contextualism. *Mind & Language* 34(2). 158–182.

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

Hansen, Mikkel B. 2010. If you know something, say something: young children's problem with false beliefs. *Frontiers in Psychology* 1.

Hansen, Mikkel B, Esben N Petersen, Arne Poulsen & Edith Salès-Wuillemin. 2017. When belief ascriptions are about more than what is on someone else's mind. *Discourse Processes* 54(8). 655–669.

Hansen, Nat. 2020. "nobody would really talk that way!": the critical project in contemporary ordinary language philosophy. *Synthese* 197(6). 2433–2464.

Hansen, Nat & Emmanuel Chemla. 2013. Experimenting on contextualism. *Mind & Language* 28(3). 286–321.

Hansen, Nat & Emmanuel Chemla. 2015. Linguistic experiments and ordinary language philosophy. *Ratio* 28(4). 422–445.

Harris, Paul L, Deborah T Bartz & Meredith L Rowe. 2017a. Young children communicate their ignorance and ask questions. *Proceedings of the National Academy of Sciences* 114(30). 7884–7891.

Harris, Paul L, Bei Yang & Yixin Cui. 2017b. 'I don't know': Children's early talk about knowledge. *Mind & Language* 32(3). 283–307.

Harris, Richard J. 1975. Children's comprehension of complex sentences. *Journal of Experimental Child Psychology* 19(3). 420–433.

Hart, Betty & Todd R. Risley. 1995. *Meaningful differences in the everyday experience of young American children*. Baltimore: Paul H. Brookes Publishing Co.

Haslanger, Sally. 2000. Gender and race:(what) are they?(what) do we want them to be? *Noûs* 34(1). 31–55.

Helming, Katharina A, Brent Strickland & Pierre Jacob. 2016. Solving the puzzle about early belief-ascription. *Mind & Language* 31(4). 438–469.

Hintikka, Jaakko. 1975. Different constructions in terms of the basic epistemological verbs: A survey of some problems and proposals. *The Intentions of Intensionality* 1–25.

Hoff, Erika. 2006. How social contexts support and shape language development. *Developmental Review* 26(1). 55–88.

Hoff, Erika & Letitia Naigles. 2002. How children use input to acquire a lexicon. *Child Development* 73(2). 418–433.

Hoff-Ginsberg, Erika Christine. 1981. *The role of linguistic experience in the child's acquisition of syntax*. University of Michigan.

Hooper, Joan B. 1975. On assertive predicates. In *Syntax and semantics*, vol. 4, 91–124. Brill.

Hopmann, Marita R & Michael P Maratsos. 1978. A developmental study of factivity and negation in complex syntax. *Journal of Child Language* 5(02). 295–309.

Howard, Alice Ann, Lara Mayeux & Letitia R Naigles. 2008. Conversational correlates of children's acquisition of mental verbs and a theory of mind. *First Language* 28(4). 375–402.

Husband, E, A Werkmann Horvat & A Gagliardi. 2016. Indexicals in shifty contexts: Problems for language acquisition. In *Proceedings of the 40th annual boston university conference on language development*, vol. 40, Cascadilla Press.

Huttenlocher, Janellen, Heidi Waterfall, Marina Vasilyeva, Jack Vevea & Larry V Hedges. 2010. Sources of variability in children's language growth. *Cognitive Psychology* 61(4). 343–365.

Ichikawa, Jonathan & C. S. I. Jenkins. 2017. On putting knowledge 'first'. In Joseph Adam Carter, Emma C. Gordon & Benjamin Jarvis (eds.), *Knowledge first: Approaches in epistemology and mind*, Oxford University Press.

Jenkins, Jennifer M, Sheri L Turrell, Yuiko Kogushi, Susan Lollis & Hildy S Ross. 2003. A longitudinal investigation of the dynamics of mental state talk in families. *Child Development* 74(3). 905–920.

Johnson, Carl Nils & Michael P Maratsos. 1977. Early comprehension of mental verbs: Think and know. *Child Development* 1743–1747.

Johnson, Carl Nils & Henry M Wellman. 1980. Children's developing understanding of mental verbs: Remember, know, and guess. *Child Development* 1095–1102.

Kaminski, Juliane, Josep Call & Michael Tomasello. 2008. Chimpanzees know what others know, but not what they believe. *Cognition* 109(2). 224–234.

Kaplan, David. 1989. Afterthoughts. In J. Almog, J. Perry & H. Wettstein (eds.), *Themes from kaplan*, 565–614. Oxford University Press.

Karttunen, Lauri. 1971. Some observations on factivity. *Research on Language & Social Interaction* 4(1). 55–69.

Kennedy, C & J Stanley. 2009. On 'average'. *Mind* 118(471). 583–646.

Kennedy, Christopher. 1999. *Projecting the adjective: The syntax and semantics of gradability and comparison*. Routledge.

Kim, Minsun & Yuan Yuan. 2015. No cross-cultural differences in the gettier car case intuition: A replication study of weinberg et al. 2001. *Episteme* 12(3). 355–361.

Kim, Sunae, Markus Paulus, Beate Sodian & Joelle Proust. 2016. Young children's sensitivity to their own ignorance in informing others. *PLoS ONE* 11(3). e0152595.

Kiparsky, Paul & Carol Kiparsky. 1970. Fact. In Manfred Bierwisch & Karl Erich Heidolph (eds.), *Progress in Linguistics: A Collection of Papers*, 143–173. The Hague: Mouton.

Kornblith, Hilary. 2002. *Knowledge and its place in nature*. Oxford University Press.

Kovács, Ágnes Melinda, Tibor Tauzin, Ernő Téglás, György Gergely & Gergely Csibra. 2014. Pointing as epistemic request: 12-month-olds point to receive new information. *Infancy* 19(6). 543–557.

Kovács, Ágnes Melinda, Ernő Téglás & Gergely Csibra. 2021. Can infants adopt underspecified contents into attributed beliefs? representational prerequisites of theory of mind. *Cognition* 213. 104640.

Kuczaj, Stan A. 1978. Children's judgments of grammatical and ungrammatical irregular past-tense verbs. *Child Development* 319–326.

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

Léger, Catherine. 2008. The acquisition of two types of factive complements. In *Language acquisition and development: Proceedings of gala 2007*, 337–347.

Lehrer, Keith & Stewart Cohen. 1983. Justification, truth, and coherence. *Synthese* 55(2). 191–207.

Leslie, Alan M, Ori Friedman & Tim P German. 2004. 'theory of mind'as a mechanism of selective attention. *Trends in Cognitive Sciences* 12(8). 528–533.

Lewis, David. 1975. Languages and Language. In K Gunderson (ed.), *Language, mind, and knowledge*, University of Minnesota Press.

Lewis, David. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74(4). 549–567.

Lewis, Davis. 1980. Index, context, and content. In Stig Kanger & Sven Ohman (eds.), *Philosophy and grammar*, 79–100. Dordrecht: Reidel.

Lewis, Shevaun, Valentine Hacquard & Jeffrey Lidz. 2017. "Think" pragmatically: Children's interpretation of belief reports. *Language Learning and Development* 1–23.

Li, Jincai. 2022. *The referential mechanism of proper names: Cross-cultural investigations into referential intuitions*. Taylor & Francis.

Lidz, Jeffrey & Annie Gagliardi. 2015. How nature meets nurture: Universal grammar and statistical learning. *Annual Review of Linguistics* 1(1). 333–353.

Ludlow, Peter. 2003. Referential Semantics for I-Languages? In Louise M. Antony & Norbert Hornstein (eds.), *Chomsky and His Critics*, 140–161. Wiley Online Library.

Machery, Edouard, Ron Mallon, Shaun Nichols & Stephen P Stich. 2004. Semantics, cross-cultural style. *Cognition* 92(3). B1–B12.

Machery, Edouard, Stephen Stich, David Rose, Amita Chatterjee, Kaori Karasawa, Noel Struchiner, Smita Sirker, Naoki Usui & Takaaki Hashimoto. 2017. Gettier across cultures. *Noûs* 51(3). 645–664.

Macnamara, John, Erica Baker & Chester L Olson. 1976. Four-year-olds' understanding of" pretend"," forget", and" know": Evidence for propositional operations. *Child Development* 62–70.

MacWhinney, Brian. 2000. *The CHILDES Project: Tools for analyzing talk. transcription format and programs*. Mahwah, NJ: Lawrence Erlbaum Associates 3rd edn.

Maffie, James. 1990. Recent work on naturalizing epistemology. *American Philosophical Quarterly* 27. 281–293.

Marinis, Theo. 2016. Acquiring Possessives. In Jeffrey Lidz, William Snyder & Joe Pater (eds.), *The Oxford Handbook of Developmental Linguistics*, Oxford University Press.

de Marneffe, Marie-Catherine, Mandy Simons & Judith Tonhauser. 2019. The commitmentbank: Investigating projection in naturally occurring discourse. In *Proceedings of sinn und bedeutung*, vol. 23 2, 107–124.

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

Mehrani, Mehdi B. 2011. What is biased? children's strategies or the structure of yes/no questions? *First Language* 31(2). 214–221.

Moll, Henrike, Sarah Kane & Luke McGowan. 2016. Three-year-olds express suspense when an agent approaches a scene with a false belief. *Developmental Science* 19(2). 208–220.

Moore, Chris, Dana Bryant & David Furrow. 1989. Mental terms and the development of certainty. *Child Development* 167–171.

Moore, Chris & Jane Davidge. 1989. The development of mental terms: Pragmatics or semantics? *Journal of Child Language* 16(03). 633–641.

Nagel, Jennifer. 2013. Knowledge as a Mental State. *Oxford Studies in Epistemology* 4. 275–310.

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

O'Neill, Daniela K. 1996. Two-year-old children's sensitivity to a parent's knowledge state when making requests. *Child Development* 67(2). 659–677.

O'Neill, Daniela K, Janet Wilde Astington & John H Flavell. 1992. Young children's understanding of the role that sensory experiences play in knowledge acquisition. *Child Development* 63(2). 474–490.

O'Neill, Daniela K & Selena Chong. 2001. Preschool children's difficulty understanding the types of information obtained through the five senses. *Child Development* 72(3). 803–815.

O'Neill, Daniela K & Alison Gopnik. 1991. Young children's ability to identify the sources of their beliefs. *Developmental Psychology* 27(3). 390.

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

Ostler, Nicholas & B.T. Atkins. 1991. Predictable meaning shift: Some linguistic properties of lexical implication rules. In J. Pustejovsky & S. Bergler (eds.), *Lexical semantics and knowledge representation*, Springer-Verlag.

Papafragou, Anna, Kimberly Cassidy & Lila Gleitman. 2007a. When we think about thinking: The acquisition of belief verbs. *Cognition* 105(1). 125–165.

Papafragou, Anna, Peggy Li, Youngon Choi & Chung-hye Han. 2007b. Evidentiality in language and cognition. *Cognition* 103(2). 253–299.

Perner, Josef, Susan R Leekam & Heinz Wimmer. 1987. Three-year-olds' difficulty with false belief: The case for a conceptual deficit. *British Journal of Developmental Psychology* 5(2). 125–137.

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

Phillips, Jonathan, Joshua Knobe, Brent Strickland, Pauline Armary & Fiery Cushman. 2018. Evidence for evaluations of knowledge prior to belief. In *Proceedings of the fortieth annual conference of the cognitive science society*, .

Pietroski, Paul M. 2005. Meaning Before Truth. In Georg Peter Gerhard Preyer (ed.), *Contextualism in philosophy*, 255–302. Oxford University Press.

Pietroski, Paul M. 2010. Concepts, Meaning and Truth: First Nature, Second Nature and Hard Work. *Mind & Language* 24(3). 247–278.

Pietroski, Paul M. 2018. *Conjoining Meanings: Semantics Without Truth-Values*. Oxford University Press.

Pillow, Bradford H. 1989. Early understanding of perception as a source of knowledge. *Journal of Experimental Child Psychology* 47(1). 116–129.

Pratt, Chris & Peter Bryant. 1990. Young children understand that looking leads to knowing (so long as they are looking into a single barrel). *Child Development* 61(4). 973–982.

Premack, David & Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences* 1(04). 515–526.

Pustejovsky, J. 1995. *The generative lexicon*. MIT Press.

Quilty-Dunn, Jake. 2021. Polysemy and Thought: Toward a Generative Theory of Concepts. *Mind & Language* 36(1). 158–185. http://dx.doi.org/10.1111/mila.12328.

Quine, W.V.O. 1960. *Word and object*. MIT Press.

Roberts, Thomas de Haven. 2021. *How to make believe: Inquisitivity, veridicality, and evidentiality in belief reports*: UC Santa Cruz dissertation.

Romero, Maribel & Chung-hye Han. 2004. On negative yes/no questions. *Linguistics and Philosophy* 27(5). 609–658.

Rothschild, Daniel & Gabriel Segal. 2009. Indexical Predicates. *Mind & Language* 24(4). 467–493.

Rubio-Fernández, Paula & Bart Geurts. 2013. How to pass the false-belief task before your fourth birthday. *Psychological Science* 24(1). 27–33.

Salter, Gideon & Richard Breheny. 2019. Removing shared information improves 3-and 4-year-olds' performance on a change-of-location explicit false belief task. *Journal of Experimental Child Psychology* 187. 104665.

Schlenker, Philippe. 2021. Triggering presuppositions. *Glossa: a journal of general linguistics (2021-...)* 6(1).

Schulz, Petra. 2003. *Factivity: Its nature and acquisition*, vol. 480. Walter de Gruyter.

Scoville, Richard P & Alice M Gordon. 1980. Children's understanding of factive presuppositions: An experiment and a review. *Journal of Child Language* 7(02). 381–399.

Searle, John R. 1975. *Indirect speech acts*. na.

Semler, Jen & Paul Henne. 2019. Recent experimental work on "ought" implies "can". *Philosophy Compass* 14(9). e12619.

Shatz, Marilyn, Henry M Wellman & Sharon Silber. 1983. The acquisition of mental verbs: A systematic investigation of the first reference to mental state. *Cognition* 14(3). 301–321.

Shope, Robert K. 1983. *The analysis of knowing. a decade of research*. Princeton: Princeton University Press.

Siegal, Michael & Karen Beattie. 1991. Where to look first for children's knowledge of false beliefs. *Cognition* 38(1). 1–12.

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

Slaughter, Virginia, Candida C Peterson & Emily Mackintosh. 2007. Mind what mother says: Narrative input and theory of mind in typical children and those on the autism spectrum. *Child Development* 78(3). 839–858.

Sodian, Beate, Claudia Thoermer & Nadine Dietrich. 2006. Two-to four-year-old children's differentiation of knowing and guessing in a non-verbal task. *European Journal of Developmental Psychology* 3(3). 222–237.

Southgate, Victoria & Angelina Vernetti. 2014. Belief-based action prediction in preverbal infants. *Cognition* 130(1). 1–10.

Sowalsky, Emily, Valentine Hacquard & Tom Roeper. 2009. Is pp opacity on the path to false belief. In *Generative approaches to language acquisition north america (galana)*, vol. 3, 263–261.

Spector, Benjamin & Paul Egré. 2015. A uniform semantics for embedded interrogatives: An answer, not necessarily the answer. *Synthese* 192(6). 1729–1784.

Spenader, Jennifer. 2003. Factive presuppositions, accommodation and information structure. *Journal of Logic, Language and Information* 12(3). 351–368.

Sprouse, Jon & Diogo Almeida. 2013. The empirical status of data in syntax: A reply to gibson and fedorenko. *Language and Cognitive Processes* 28(3). 222–228.

Sprouse, Jon, Carson T Schütze & Diogo Almeida. 2013. A comparison of informal and formal acceptability judgments using a random sample from linguistic inquiry 2001–2010. *Lingua* 134. 219–248.

Starmans, Christina & Ori Friedman. 2020. Expert or esoteric? Philosophers attribute knowledge differently than all other academics. *Cognitive Science* 44(7). e12850.

Strawson, Peter F. 1950. On referring. *Mind* 320–344.

Tardif, Twila & Henry M Wellman. 2000. Acquisition of mental state language in mandarin-and cantonese-speaking children. *Developmental Psychology* 36(1). 25.

Taylor, Marjorie, Bonnie M Esbensen & Robert T Bennett. 1994. Children's understanding of knowledge acquisition: The tendency for children to report that they have always known what they have just learned. *Child Development* 65(6). 1581–1604.

Thomas, Ruthann C & Larry L Jacoby. 2013. Diminishing adult egocentrism when estimating what others know. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 39(2). 473.

Turri, John. 2013. A Conspicuous Art: Putting Gettier to the test. *Philosophers' Imprint* 13(10). 1–16.

Uegaki, Wataru. 2015. *Interpreting questions under attitudes*: Massachusetts Institute of Technology dissertation.

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

Valente, Matheus. 2019. Communicating and disagreeing with distinct concepts: a defense of semantic internalism. *Theoria* 85(4). 312–336.

Ventura, Rafael. 2019. Quantitative methods in philosophy of language. *Philosophy Compass* 14(7). e12609.

Vicente, Agustín. 2021. Approaches to Co-predication: Inherent Polysemy and Metaphysical Relations. *Journal of Pragmatics* 182. 348–357.

de Villiers, Jill G. 1995. Questioning minds and answering machines. In *Proceedings of the 19th boston university conference on language development*, 20–36.

Vogel, C. A. 2020. Realism and Lexical Flexibility. *Theoria* 86(2). 145–186.

Vogel, Christopher A. 2016. Lexical Flexibility, Natural Language and Ontology. *Croatian Journal of Philosophy* 16(1).

Vogel, Christopher A. 2018. Externalism and Conceptual Analysis. *Metaphilosophy* 49(5). 730–65.

Weinberg, Jonathan M., Shaun Nichols & Stephen Stich. 2001. Normativity and Epistemic Intuitions. *Philosophical Topics,* 29(1-2). 429–460. http://dx.doi.org/10.5840/philtopics2001291/217.

Weist, Richard M, Hanna Wysocka & Paula Lyytinen. 1991. A cross-linguistic perspective on the development of temporal systems. *Journal of Child Language* 18(1). 67–92.

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

Westra, Evan & Peter Carruthers. 2017. Pragmatic development explains the theory-of-mind scale. *Cognition* 158. 165–176.

White, Aaron Steven. 2021. On believing and hoping whether. *Semantics and Pragmatics* 14. 6.

White, Aaron Steven & Kyle Rawlins. 2018. The role of veridicality and factivity in clause selection. In *Proceedings of the 48th annual meeting of the north east linguistic society*, 221–234.

Williams, Michael. 2001. *Problems of Knowledge*. Oxford University Press.

Williams, SG. 2002. Ambiguity and belief1. *Royal Institute of Philosophy Supplements* 51. 253–278.

Williamson, Timothy. 1994. *Vagueness*. Routledge.

Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford University Press.

Wimmer, Heinz, G-Jürgen Hogrefe & Josef Perner. 1988. Children's understanding of informational access as source of knowledge. *Child Development* 386–396.

Woolley, Jacqueline D & Marc J Bruell. 1996. Young children's awareness of the origins of their mental representations. *Developmental Psychology* 32(2). 335.

Woolley, Jacqueline D & Henry M Wellman. 1993. Origin and truth: Young children's understanding of imaginary mental representations. *Child development*

64(1). 1–17.

Zagzebski, Linda. 1994. The Inescapability of Gettier Problems. *Philosophical Quarterly* 44(174). 65–73.

Zalta, Edward N. 2001. Fregean senses, modes of presentation, and concepts. *Philosophical Perspectives* 15. 335–359.