The development of negation in language and thought

Annika McDermott-Hinman* & Roman Feiman

Department of Cognitive, Linguistic, and Psychological Sciences, Brown University

* Corresponding author. Address: Room 241, 190 Thayer St., Providence, RI, USA, 02906. Email: annika_mcdermott-hinman@brown.edu

1 Introduction

By using linguistic expressions of negation, people express *negated thoughts*: thoughts that something is *not* the case, that a predicate does *not* apply to a subject, that one kind of thing is *not* another, and so on. If being able to express a thought requires being able to think it in the first place, then expressing negation is not one ability, but two. It requires first, the capacity to think a thought that includes the concept NEGATION, and second, a mapping between this NEGATION concept and the word(s) that express it in one's language.

If these are two abilities, there are correspondingly two separable questions about their development. First, how does the concept NEGATION emerge in thought over both phylogenesis and ontogenesis? Second, how do children learn to map this concept onto the particular construction(s) that express it in their language?

1.1 The target of development

To consider each of these problems, it helps to be as precise as we can about what the end state of development looks like. Unfortunately, given that there are many different ways to define negation in different logical systems, as well as many different theories of what expressions of negation mean in natural languages (see Horn, 1989), precision requires making choices that may not generalize to all logical formulations or linguistic theories. We define our working concept NEGATION as the truth-functional propositional operator that reverses the truth value of the proposition with which it composes: if a proposition p is true, then the NEGATION of p must be false, and vice versa. Note that this definition of the concept has it performing the same role independent of the content of p. The word *not* has the same meaning—NEGATION—in the expressions *the ball is not red* and *the toy is not in the bucket* and *Pete Buttigieg is not the mayor of New York.* While standard in definitions of negation, this feature will become important when we consider conceptual precursors to negation in Section 5.

Although we will focus on sentential negation as a logical operator defined over propositions here, our commitment is meant to be minimal. We do not make any strong claim that sentential negation is the concept NEGATION that people use to think negated thoughts and that constitutes the meaning of negative expressions (as opposed to, e.g., set theoretic complementation). We do not even want to commit to there being any single such concept. Rather, we take this definition to be useful as a testbed for considering how a particular version

of the concept and its expression could develop. We hope that the reader might then, mutatis mutandis, profitably apply the framework we develop here to alternative formulations and theories of negation.

The benefit of being explicit about the target of development is that it makes it clear that possession of this concept imposes stiff requirements on the kinds of representations that must be in the mind more generally. If NEGATION is a mental computation articulated over propositions and truth-values, then minds that can think negated thoughts must correspondingly be minds that can, on some level of description, represent propositions and truth-values. The force of this requirement will become evident in Section 5, where we consider how a concept of negation could possibly develop.

1.2 Overview

This chapter aims to showcase the recent advances that have been made by considering the two questions of conceptual development and language acquisition separately. While being able to express negation requires having both the concept and the words for it, conceptual development and language acquisition need not be equally hard problems for the learner. Focusing on English, where there is currently the most evidence, Section 2 describes the developmental trajectory of children's word-to-concept mapping, taking it as the explanandum. Section 3 then shows how recent studies have disentangled the relative contributions of conceptual development and language acquisition to this trajectory.

Section 4 focuses on the problem of language acquisition. Since no child comes into the world knowing how their particular language expresses negation, the general problem is the same for all children—they must figure out which of the words in their input map to NEGATION. However, languages might differ in how hard they make this problem in several ways. Section 4 focuses on structural differences in how easy a language makes it to rule out competing hypotheses about the meanings of negative expressions.

While language acquisition is necessarily a common problem for all children, whether conceptual development is a problem for any child depends on how much of the developmental process has been completed by evolution prior to birth. If humans are innately endowed with the capacity to think negated thoughts, then the mechanisms of conceptual development would be located in phylogenesis rather than ontogenesis, and uncovering them would require studies in

comparative cognition rather than developmental psychology. While evidence from phylogenesis is beyond the scope of this chapter, Section 5 highlights the way in which the very same questions about the mechanisms of the conceptual development of NEGATION arise whether this development is located in evolution or in infancy, and reviews the latest evidence of conceptual development in ontogenesis.

2 Children's negative expressions

Listening to the speech of toddlers, it is easy to wonder whether they've missed the theorist's memo that expressions of negation should be hard to learn. Across languages, negators are often among children's first words. In English, for example, two thirds of parents report that their children are saying *no* by 17 months, an age at which the median child produces only about 40 words (Fenson et al., 2007; Frank et al., 2017). Production, however, might overestimate competence. Does the one-year-old screaming "No!" in the supermarket really mean to express the truth-functional concept, NEGATION, or could they instead be communicating something simpler—maybe a feeling of dislike or disappointment?

In support of this leaner possibility, careful studies of children's speech and comprehension have gradually revealed that early uses of negation are lexically, syntactically, and semantically limited. Semantically, children's early negative productions appear to serve specific communicative functions, expressing many fewer meanings than the full capability of the truth-functional concept NEGATION. For example, Bloom (1970) argued that children's earliest uses of *no* were not negating the truth of any proposition, but were rather being used for the more specific communicative function of rejecting or opposing something that the child was offered, asked, or instructed to do. Under Bloom's account, using *no* only to express rejection is the first of multiple stages through which children typically progress on the way to mastering logical negation. During this progression, Bloom argued, children not only learn new negative expressions, they also acquire new negative meanings to express. After rejection, which children initially express using single-word utterances of *no*, Bloom argued that children begin to use multi-word negative utterances to comment on nonexistence (something being absent where it was expected to be present), followed by multi-word expressions of rejection. Only later, around their second birthday, did the children Bloom studied finally begin to deny the truth of what someone else had said—the clearest expression of the truth-functional concept NEGATION.

2.1 Limited uses reflect limited meanings

Since Bloom's seminal investigation of children's speech, there has been much more consensus that children initially use negative expressions in more limited ways than adults than there has been consensus about what exactly those limited uses are. Expanding on Bloom's proposed functions of rejection and nonexistence, Choi (1988) agreed that children are late to use negation words for denial but argued on the basis of children's productions in English, French, and Korean that earlier negative utterances express not only rejection and nonexistence, but also distinctly failure, and that along with the emergence of denial, they also start to express inability, epistemic negation, normative negation, and inferential negation (see also Cameron-Faulkner, Lieven, & Theakston, 2007; Nordmeyer & Frank, 2018). De Villiers & De Villiers (1979) agreed that English learners' early negative expressions had the distinct meanings of rejection and nonexistence, but also posited three other categories: disappearance, nonoccurrence, and cessation. Gopnik and Meltzoff (1985) dropped nonexistence from the list and suggested that children initially use *no* for refusal, then later to express the failure of a plan, and finally in a third stage to mean propositional NEGATION. Thus, while there is general consensus that the earliest uses of negation express rejection or refusal, there is not consensus about which intermediate meanings children's negative utterances express and in what order they appear. What is crucial, however, is that in each of these proposals, children's early uses of negative expressions are limited. Instead of expressing the full variety of meanings that would be expressible if they took words like *no* and *not* to mean truth-functional NEGATION, children initially use such words for a small number of more specific communicative functions.

Bloom (1970) suggested that changing uses reflect changes in the concepts or meanings that children assign to negative expressions, with the concepts REJECTION and NONEXISTENCE being distinct from the truth-functional NEGATION concept that underlies denial uses (see also Pea, 1980b). On the other classification schemes, each use could likewise be a candidate different concept.

However, the inference from how someone tends to use a word to what they take that word to mean is not straightforward. For example, it is possible for a word to be used for rejecting something without having the concept REJECTION as a part of its semantics. Consider the statement *I do not want a cookie* in response to being offered a cookie. Although this

statement can be used to reject a cookie, and although the word *not* is integral to the formation of that meaning, *not* itself does not mean REJECTION—it means NEGATION, and the communicative function of rejecting only results from *not* composing with the predicate *want a cookie* in the context of the preceding cookie offering. Although young children's negation utterances are syntactically simpler than adults', it is possible that they are expressing NEGATION from the start. The trajectory of changing functions might, in that case, be due to younger children simply having more occasion to reject things ("Don't want!") and complain about the disappointing nonexistence of something they did want ("No more!") than they have to comment on the falsity of propositions relative to the world. Worse yet, maybe the changing functions are an artifact of the adult theorist's interpretation of the intent of the child who is limited to producing only one or two words at a time. It is possible, for example, that it is simpler to understand what a child means when they are rejecting with a cry of "No!" in response to an offer of something they obviously do not want, than when they are using the same single-word *no* to deny that a prior statement is true. In either of these cases, the observed developmental trajectory would be an artifact—either of children's most often and earliest employed communicative functions or of the theorist's ability to interpret child speech—and would not reflect changes in the meanings of their negative expressions.

Two sources of evidence suggest that children's limited uses do result from their mapping of several concepts other than NEGATION to the forms of negative expressions. First is that related but distinct negative concepts seem to exist in adulthood, as evidenced by the fact that many of the world's languages lexicalize these concepts as distinct words or constructions. Reviewing 95 languages, Veselinova (2013) found that 42 use entirely different forms to express nonexistence than truth-functional NEGATION. Though we are not aware of a systematic review, several other languages assign different forms to other functions, as well. For example, in Tagalog, in addition to there being a word expressing NEGATION and another expressing nonexistence, there are also specific words to communicate the meanings of rejection, prohibition, and epistemic negation (Schachter & Otanes, 1983). Korean also has distinct forms for each of these functions, and an additional form for inability. Moreover, children learning Korean begin to produce the form for NEGATION later than the forms for nonexistence or rejection (Choi, 1988), following the sequence Bloom (1970) had first found for different functions of *no* in English. That the meanings of rejection and nonexistence are lexicalized

separately in unrelated languages suggests that what Bloom, Choi, and others identified as different functions of negation might be represented as distinct concepts.

Additional evidence that limited functions express limited concepts comes from tests of children's comprehension of negative expressions. In laboratory tasks, children start to understand that *no* expresses NEGATION more than half a year later than their first uses of the word. In one study, children were given the opportunity to search for a toy in one of two locations based on a verbal clue. When given positive clues (e.g., "It's in the bucket"), the youngest age group that was tested (20-month-olds) succeeded. However, when the clue was negative, regardless of whether the negator used was *not* (e.g., "It's **not** in the bucket") or *no* (e.g., "Is it in the bucket?" "**No**, it's not"), children did not succeed robustly until 27 months of age (Feiman, Mody, & Carey, 2017). Virtually all of the children in this study were already producing the word *no* by this point, according to their parents' report. In another study, using a nearly identical method but different exclusion criteria, Austin, Theakston, Lieven, & Tomasello (2014) found success with both *no* and *not* at 24 months. The comprehension of logical negation around the start of the third year of life converges with the age at which studies of children's production find them starting to use negative expressions in ways that clearly express a concept of truth-functional negation (Bloom, 1970; Cameron-Faulkner, et al., 2007; Choi, 1988; Gopnik & Meltzoff, 1985; Nordmeyer & Frank, 2018; Pea, 1980a). Thus, these comprehension studies provide additional evidence for the hypothesis that the limited functions for which young children use negative expressions reflect more limited concepts (e.g., NONEXISTENCE and REJECTION). If 17-month-old children had already mapped *no* to NEGATION, they should have understood *no* to mean NEGATION in these comprehension tasks at the same age at which they had first produced it themselves.

2.2 A single bottleneck for learning that both *no* and *not* express NEGATION

Strikingly, though the age of success varied slightly between the studies, both Austin, et al. (2014) and Feiman, et al., (2017) found that children understood the two words *no* and *not* as expressing NEGATION at the same age within each study. This, despite the fact that children begin to produce *no* much earlier than *not* and despite the differences in the two words' syntactic distributions and compositional properties. As Feiman, et al. (2017) argued, this pattern suggests some common bottleneck to acquiring the NEGATION meaning of both words that transcends their

differences. Feiman, et al. (2017) go on to identify two possible, distinct causes of this bottleneck: conceptual development (learning to think the concept NEGATION), and language acquisition (learning the mapping between NEGATION and both word forms, *no* and *not*).

On the first alternative, children must undergo some further conceptual development between 17 months—the age at which most children are producing *no*—and 24 months—the age at which Austin et al. (2014) found children to comprehend truth-functional *no* and *not*—in order to be able to think negated thoughts, independent of learning how those thoughts can be expressed in their language. Perhaps children initially only have access to related, but more limited concepts like REJECTION and NONEXISTENCE, and so they initially hypothesize that those concepts are the meanings for negative expressions, since they are consistent with some (even if not all) adult uses of those expressions. Under this account, children must eventually acquire the concept NEGATION, either by constructing it out of other concepts, or by an independent maturational process that occurs at some point after they first begin to say *no*. Only once they gain the ability to think negated thoughts would children then be able to learn the adult meanings of the words *no* and *not,* mapping them to NEGATION.

Alternatively, it is also possible that the common bottleneck on the acquisition of this mapping is due to the challenges of language learning alone. Even if young children have no trouble thinking negated thoughts, they must still figure out how their language expresses NEGATION. Expressions of negation are prime examples of what Gleitman and colleagues (2005) called "hard words" to learn—words that cannot be associated with anything reliably co-occurring with their use in the child's environment. The chief role of negative expressions is to systematically modify the meanings of other content with which they compose, which could itself vary widely. If learning that *no* and *not* map to NEGATION requires noticing how these words modify the meaning of the rest of the utterance in which they appear, understanding the rest of that utterance would probably be helpful. Note that in this case, the learning problem facing children would be just the same as the problem facing any conceptually developed adult who is just starting to learn a new language. Imagine being shown a blue ball and hearing the Tagalog expression, *hindi pula ang bola*. How hard is it to understand what *hindi* means, if you do not already know what *pula, ang,* and *bola* mean? How much easier does the task of understanding *hindi* become once you learn that *ang bola* means THE BALL and *pula* means RED? Similarly, young English-learning toddlers who hear, *the ball is not red*, may not be able

to understand enough of that sentence to guess that *not* means NEGATION, or that *no* means the same in other instances, even if they can get enough evidence to guess that *no* means REJECTION, and even if they are perfectly capable of thinking with both concepts and entertaining both as hypotheses. Language learners might only be able to map both *no* and *not* to NEGATION once they have learned enough of the rest of their language to guess that this is the concept that those words express.

3 Linguistic information as a bottleneck

Recent studies have tested whether insufficient conceptual development or insufficient linguistic information might be the cause of the bottleneck that English-learning children face in mapping NEGATION to both *no* and *not.* In particular, if extracting enough linguistic information from the input is the main bottleneck, learners who are more conceptually sophisticated and cognitively developed than 1-year-old toddlers should find it similarly hard to acquire these expressions if their linguistic input is similarly limited.

3.1 Evidence from the Human Simulation Paradigm

In a recent study, Gomes and colleagues (2021; in prep) employed a version of the Human Simulation Paradigm (HSP) to test this prediction. Originally proposed by Gleitman (1990) and first used by Gillette, Gleitman, Gleitman, and Lederer (1998) to tease apart informational and conceptual factors in language acquisition, the HSP is a method that simulates different kinds of informational restrictions imposed on the language learning environment of otherwise conceptually mature adults.

In Gomes and colleagues' experiments, adults were presented with videos of parent-child interactions in which all of the speech had been silenced, and were then asked to guess what a parent had said at the moment in the video when a beep had sounded. Participants in one condition were presented with a video in which all the rest of the speech other than the beeped-out word was subtitled, and therefore saw the full linguistic context surrounding the target they were asked to guess, while in another condition participants were provided with the silent video only, and no other information. The Human Simulation Paradigm tests how manipulating the amount of linguistic information that is available to a learner—while lifting any limit that might be due to their conceptual development—affects their ability to tell that NEGATION is being

expressed. The adult participants in this study had access to the concept NEGATION and even knew the words to express it. Their task was only to identify when a word was being used to express that particular concept.

The video-only condition simulates the linguistic information available to preverbal infants in the process of trying to learn their first words—they may know that words are being spoken but must rely on non-linguistic perceptual information to guess their meanings. Recall that at this point in acquisition, children only say *no* and use it only for rejection. The video-with-language condition simulates the linguistic information available to a later learner who has already learned much of their language and can use the remainder of the utterance to infer the missing word. At this point in children's language acquisition, they start to use *not* as well as *no*, and use both words to express truth-functional negation.

Adult participants' performance on this HSP task mirrored children's changing uses of negative expressions as they learn more language and acquire more linguistic information. In the video-only condition, when the target word was *no*, participants guessed this word correctly over half of the time when parents had used it to express prohibition, suggesting that the visual scene alone contained enough information to figure out that a parent was prohibiting a child. In contrast, participants guessed *no* correctly only 10% of the time when parents had used it for what Bloom called 'denial', an expression of truth-functional NEGATION that cannot be understood using one of the narrower negative concepts (see Pea, 1980b). When the target word was *not*, which parents tend to use more often for denial and less often for prohibition (Cameron-Faulkner, et al., 2007), none of the participants guessed it correctly in the video-only condition no matter which way it was used. In contrast, in the condition where linguistic context was provided in addition to video, participants' pattern of correct guesses was very different. Their ability to guess both words when they expressed NEGATION was greatly improved: over 50% of their guesses were correct, both when denial uses of *no* and *not* were the targets. It appears that linguistic context was necessary for adults to be able to guess that the concept NEGATION was being communicated.

The results of this study also help to explain children's earlier, more limited uses of *no*. When participants saw *no* being used to prohibit the child, they were usually able to guess the meaning correctly with no additional linguistic context. They were not able to do the same for *not*. Not only was prohibition the only use of *no* that allowed participants to reliably guess the

word in the video-only condition, participants were actually more accurate at guessing this use of *no* than they were at guessing even common, early acquired concrete nouns (e.g., *dog* and *ball*). This suggests that prohibition uses of *no,* unlike NEGATION uses of *no* or prohibition uses of *not,* do have reliable perceptual correlates. As Choi (1988) suggested, adults' prohibitions might reflect the same meaning of *no* as children's rejections. In both uses, the word *no* might express that the speaker feels negatively about something or is requesting that something stop happening. Spoken by a parent to a child, it tends to address a child's undesirable behavior, but spoken by a child to a parent, it tends to address a parental command or offer. The findings from Gomes and colleagues' experiments suggest that children might learn that *no* expresses this concept first because this is the meaning that is easiest to guess in the absence of richer linguistic knowledge.

In sum, adults' pattern of successful guessing under different conditions of information availability mirrors the trajectory of children's acquisition. Early in the acquisition process, children produce mainly *no*, using it largely to express the communicative functions of rejection and prohibition (e.g., Bloom, 1970; Choi, 1988; Nordmeyer & Frank, 2018). Later in acquisition, once they have learned more about their language, children begin to employ the word *no* for other functions, including nonexistence and truth-functional negation, and they also begin to use the word *not* to express all of these functions.

This suggests that a linguistic information limit is sufficient to explain why children learn that *no* and *not* express NEGATION later than they start to say *no*. As they learn more of their language, the linguistic information available to them grows richer, allowing them to home in on the hypothesis—possibly previously available to thought, but not easy to find support for—that *no* and *not* express NEGATION.

3.2 Evidence from international adoption

Converging evidence for this conclusion comes from research on how negative expressions are acquired by internationally adopted children, who provide a natural experiment that disentangles cognitive development from language acquisition (Snedeker, Geren, & Shafto, 2007; 2012). When children who are born into a different language environment are subsequently adopted internationally into English-speaking homes as preschoolers, they find themselves facing a similar language learning problem to infants for whom English is their birth language. Unlike infants, however, they can bring to bear the cognitive and conceptual resources

provided by several years of additional development. Feiman, Zimmerman, and Snedeker (2019; in prep) tracked eight such children who were between the ages of 2 years, 5 months and 5 years, 6 months when they were adopted. Typically, by age two and a half, most children have already begun to understand truth-functional negation in their first language (Austin, et al., 2014; de Carvalho, et al., 2021; Feiman, et al., 2017), so most, if not all of these children had likely mapped negation to its expression in their birth language prior to their adoption. Each adopted child in the corpus was paired with a first-language learner as a control, matched on the mean length of utterance (MLU) in their speech at each sample over the study period. The number of negative utterances that were coded as having the communicative functions of rejection, nonexistence, and denial were compared between groups. If conceptual development is needed before typical toddlers can express truth-functional negation, then we would expect the adopted children, who had likely already undergone that conceptual development in acquiring their first language, to start producing more denials at a comparatively earlier point in the course of language acquisition than typical first-language learners. If, on the other hand, the main driver of children's acquisition of expressions of NEGATION is the amount of linguistic information available to the child, then we would expect adoptees and first language learners to follow similar trajectories.

Compatible with the results from the Human Simulation Paradigm and with a purely language-information bottleneck, evidence from internationally adopted children showed that more conceptually mature language learners followed a very similar acquisition trajectory to younger, first-language learners. Replicating findings from many prior studies (Bloom, 1970; Pea, 1980b; Nordmeyer & Frank, 2018), first-language learners initially produced mostly rejections, which were gradually overtaken by denials. Crucially, the adopted children followed exactly the same trajectory, and there was no significant interaction of MLU (as an index of language acquisition) by group membership (adoptee vs. first-language learner). Whether children were older internationally adopted learners or typical English-learning toddlers, their production of truth-functional negations increased at an indistinguishably similar rate as they learned more English words and the complexity of their utterances increased. Together, these two sets of findings suggest that the pattern of early rejections giving way to truth-functional uses of negative expressions over time can be fully explained by the early learners' limited

ability to extract linguistic information from their language input and to use this information to constrain their hypotheses about which words map to the concept NEGATION.

Nevertheless, evidence from both the Human Simulation Paradigm and international adoption is limited in its ability to explain how negative expressions are acquired. First, The evidence only suggests that positing conceptual development is not necessary to explain the trajectory of acquisition; neither can rule out the possibility that some conceptual development does occur in ontogenesis, which enables children to think negated thoughts and thus hypothesize that NEGATION maps onto linguistic expressions such as *no* and *not*. Second, whether or not conceptual development occurs, this evidence does not tell how children actually solve the mapping problem of language acquisition, or what information they use in that process to constrain hypothesized mappings. To address both of these questions, we turn to examine cross-linguistic evidence and consider how structural differences in the expression of negation across languages might affect acquisition and further constrain the possibilities for early conceptual development.

4 Cross-linguistic differences and the mapping problem

How might the distribution of mappings between negative concepts and negative expressions help or hinder a language learner? One cross-linguistic distinction that might impact acquisition is whether the word(s) a language uses to express negation are also used to express other distinct negative functions. This is the case, for example, in English, where the word *no* can express not only truth-functional negation, but also rejection, nonexistence, and possibly many others. If, as we argued in Section 2, these separate communicative functions reflect distinct concepts (e.g., REJECTION and NONEXISTENCE), all part of the semantics of the word *no*, then the language acquisition task that children face in learning English *no* involves mapping at least three distinct concepts onto the word. Contrast this with languages like Hebrew, Korean, or Tagalog, in which there are finer-grained mappings between words and concepts. In Tagalog, for instance, there is a word—*hindi*—that means NEGATION, and which can combine with verbs of desire or existence to convey rejection or nonexistence. But there are also additional, distinct words with meanings that can only be used for these specific functions, and that therefore directly express the concepts REJECTION (*ayaw*) and NONEXISTENCE (*wala*), among others, without having to compose with the corresponding verbs. Thus, unlike children learning English,

Tagalog-learning children must only ever learn the mapping from one negation concept to one negation word.

Which of these mapping patterns is easier to learn depends on exactly which tasks children have to solve during the process of learning negative language. We argued in Section 3 that conceptual development is not necessary to explain the common bottleneck on learning the NEGATION meaning of *no* and *not*. However, this does not rule out the possibility that such development occurs during children's acquisition of negation language. It may be the case that, at the outset of the learning process, the concept NEGATION is not initially known to children, and must be constructed out of other, earlier concepts, or precursors. Section 5 discusses potential precursors to NEGATION and how this construction process might take place, but in this section we will focus on the predictions generated by a specific proposal that is common in the negation acquisition literature: that the concept NEGATION is constructed out of the earlier negative concepts found in children's production.

In one such account, Pea (1980b) argues that children's earliest mapping of the word *no* is to the meaning PROHIBITION, and that the "affirmative-negative contrast" inherent to understanding this meaning forms the child's conceptual basis for later developing the concept truth-functional NEGATION. Another argument in this vein is that children's earliest negative expressions are metalinguistic comments (see Horn, 1989). On this view, when children initially appear to be denying the truth of an assertion, what they are really saying is *don't say that* (Drozd, 1995) or *that statement violates the rules* (e.g., of a naming game) (Hummer, Wimmer, & Antes, 1993). These metalinguistic uses might reflect a mapping of PROHIBITION or REJECTION to *no*, with the full concept NEGATION somehow developing later out of these other concepts (Dimroth, 2010; Drozd, 1995; Hummer et al., 1993). While the construction process itself is underspecified, there is some evidence to support the metalinguistic interpretation of children's early denials. Drozd (1995) found that children's apparently ungrammatical pre-sentential uses of *no* (e.g., *no the sun shining*) largely echo some previous utterance said to the child (e.g., *the sun is shining*). Here, a statement such as *no the sun shining* might be paraphrased as *don't say the sun (is) shining*, a metalinguistic rejection of the act of stating that the sun is shining, rather than a denial of the truth of the statement *the sun is shining*.

To our knowledge, there is no mechanistic account on offer of how exactly NEGATION could be constructed out of narrower concepts such as REJECTION (metalinguistic or otherwise),

NONEXISTENCE, and so on. We have trouble seeing how such an account would work because, by definition, these narrower concepts do not involve truth-values, which are the primitive elements that truth-functional negation operates over (if they did, the concept would just be NEGATION operating on a restricted domain; see also Macnamara, 1986). Nevertheless, as we discussed above, many authors have suggested such a progression in conceptual development. Even without specifying the proposal further, we think it is possible to spell out the predictions that any such progression would make, and therefore to test it without having to specify a mechanism by which it occurs.

If it is the case that the concept NEGATION is constructed from earlier-acquired negative meanings, then the cluster of related meanings mapped to a single word in a many-to-one mapping system may provide some cue to learners about which concepts to use in that construction process. Take, for example, English, where a single word—*no*—means, among other things, both REJECTION and NEGATION. Children learning English learn the REJECTION meaning early on (Bloom, 1970; Choi, 1988; Gopnik & Meltzoff, 1985). When they then notice that *no* is used in more contexts than can be accounted for by REJECTION alone, this may motivate them to construct the concept—which turns out to be NEGATION—that builds on REJECTION and encompasses these contexts. Thus, if children must construct the concept NEGATION, and if that construction relies on narrower negative meanings that get mapped to the same word as NEGATION in some languages (a many-to-one mapping) and to different words in other languages (a one-to-one mapping), then constructing the concept of NEGATION should be easier in many-to-one mapping languages than in one-to-one mapping languages.

The prediction is less clear if, on the other hand, conceptual development is not a factor, which would be consistent with evidence from the Human Simulation Paradigm and from international adoption. If NEGATION is either innate, emerges through maturation rather than construction, or else is constructed before children begin learning the language for it, then the only task facing children during negation acquisition is that of learning the mappings between negative expressions and concepts.

In that case, which kind of language is easier to learn might depend on whether children represent different narrower negative concepts as being related to each other or not. A single word expressing both REJECTION and NONEXISTENCE would be a homonym for the child who does not represent these meanings as being related through NEGATION. In that case, children may

try to avoid introducing homonyms into their lexicon and therefore may prefer to *avoid* assigning these multiple meanings to a single word (Casenhiser, 2005; Dautriche, Fibula, Fievet, & Christophe, 2018; Mazzocco, 1997; Peters & Zaidel, 1980). A one-to-one mapping, as in Tagalog or Korean, would then be easier to learn because each concept in these languages gets its own word. However, if children represent the narrower negative meanings as related, then many-to-one systems of negation like English and Spanish might not slow down the mapping process, and could even speed it up. Children have an easier time learning that new nouns apply to multiple referents when those referents share a family resemblance compared to when they appear to be unrelated (Floyd & Goldberg, 2021), though it is unclear if this should generalize to abstract concepts that do not clearly have extensions. Still, a cluster of related negative meanings on a single form might provide some advantage to the learner by helping them narrow the search space when attempting to identify new meanings for that word. A child who already has NEGATION in the hypothesis space and who already knows that the word *no* means REJECTION and NONEXISTENCE may have an easier time identifying the concept NEGATION as an additional meaning—because it is semantically related to the existing meanings—than a child who has to identify that meaning without this guidance. The fact that children may already have the concept NEGATION available to thought does not imply that they will have the concept at the tip of their proverbial tongues when considering what *no* means. They still must pick it out of the countless concepts that they are capable of thinking, and some clue as to where in the hypothesis space to search may be helpful.

Consequently, because clustered meanings would support a conceptual construction process, if a one-to-one mapping is in fact easier to learn (e.g., if Tagalog-learning children map NEGATION to *hindi* earlier in acquisition than English-learning children map NEGATION to *no*), then it is unlikely that children construct NEGATION from narrower REJECTION or NONEXISTENCE. The opposite pattern of results (earlier mapping of NEGATION to *no* in English) would be ambiguous, pending a clearer picture of whether children initially represent the narrower negative concepts as related and how this might affect the process of hypothesis-testing they engage in to map these concepts to words. Future work can test these predictions by comparing the pattern of acquisition in many-to-one languages like English and Spanish to one-to-one languages, like Tagalog and Korean.

5 Conceptual development and conceptual precursors

Whether or not conceptual change is required in ontogenesis, it is surely required in phylogenesis. Even if human infants are innately equipped with the concept of negation and only have a mapping problem left to solve, this concept must have emerged in some ancestral organism earlier in the course of evolution. Whether over phylogenetic or ontogenetic time, the concept of truth-functional negation must therefore have a developmental precursor—some mental representation that exists before NEGATION, and that plays a direct causal role in its emergence. As we discussed in Section 4, REJECTION, NONEXISTENCE, and other more limited negative concepts have been offered as candidate precursors in ontogenesis, motivated by evidence that young children map these concepts to negative expressions earlier than they map NEGATION. In contrast, research on the logical capabilities of non-human animals has focused on a different candidate precursor in phylogenesis: CONTRARIETY. In this section, we cast these candidates in the light of a framework for thinking about conceptual precursors of NEGATION more generally, and for considering precursors in tandem with the developmental process that could build NEGATION on their basis.

What should a representation be like in order to count as a precursor of NEGATION? Answering this question requires specifying: (a) in what way the precursor is sufficiently related to truth-functional negation that some process of change could construct negation on the basis of the precursor, and (b) how exactly a precursor of negation is different from negation itself. Following Feiman, Mody, & Carey (2022), we suggest that precursors of negation might differ from the full-fledged concept in two distinct ways. First, they might be internal to and encapsulated within particular mental computations, and be unavailable to perform functionally similar roles over other types of mental content (i.e., they might be *domain specific*). Second, they might be able to perform only some, but not all, of the computational role of truth-functional negation over the content with which they combine (i.e., they might only have *partial function*). Precursors with partial function might support the construction of full-fledged truth-functional negation by the addition of the missing functions of negation. Domain-specific precursors might support the same construction by the addition of domains the operators can range over, or by concatenation with other domain-specific but functionally equivalent operations, followed by a form of variable abstraction over the inputs across the domains. A precursor of negation could have one or both of these features.

Concepts like REJECTION and NONEXISTENCE are first and foremost candidate precursors in the sense of domain specificity. Their conceptual role could be functionally equivalent to the role of the complex formed by the composition of truth-functional NEGATION with a specific concept, such as WANT or EXIST. In addition, and separately from their domain specificity, these concepts could also be precursors in the sense of fulfilling only a partial function of negation. For example, the child's concept REJECTION might be not just that it is false that they want something, but a more specific feeling of aversion. While a truth-functional negation of wanting could encompass feelings ranging from aversion to indifference, aversion in particular would constitute a logically stronger polar *contrary* of WANT.

5.1 Contrariety as a precursor to negation.

Contrariety itself is the candidate precursor of negation that has received the most attention in research on the phylogenesis of negation, which has focused on the logical capacities of non-human animals. Negation is commonly defined as an operator that obeys both the Law of Non-Contradiction (LNC) and the Law of Excluded Middle (LEM) (see Horn, 1989). That is, given two propositions p and q, q is the NEGATION of p if and only if p and q cannot both be true (LNC), and either p or q must be true (LEM). Together, these properties have the consequence of defining negation as reversing the truth-value of a proposition—if p is false, then the negation of p must be true, and if p is true, its negation is false. If it is false that *you are in New Jersey*, then it must be true that *you are not in New Jersey* (at least as long as presuppositions of existence are satisfied; see Horn, 1989). Unlike negation, contrariety obeys only the Law of Non-Contradiction (LNC) and need not obey the Law of Excluded Middle (LEM). That is, two propositions are contrary to each other if they cannot both be true, but may both be false. For instance, *you are in New Jersey* is a contrary of *you are in outer space*—both of these propositions cannot be true, but both can be false if you are in some other location. Similarly, *I hate these carrots* could be a contrary to *I like these carrots*, and could in fact be what the child intends to express by a rejection use of *no.* By obeying one of the two laws that constitute negation, contrariety is a precursor of negation in the sense of performing a *partial function.* Importantly, however, nothing in the definition of contrariety necessitates that it also be *domain-specific.* In principle, any two propositions whatsoever can be evaluated for whether they are

contraries, and any proposition *p* can be represented as a contrary to any proposition *q* as long as *p* and *q* cannot both be true.

We think that prior discussions of contrariety as a potential precursor of negation in phylogenesis have blurred the distinction between partial function and domain specificity. In the literature on non-human animals' capacity for logical representations, contrariety has been assumed to be domain-specific, requiring prior experience with the incompatibility (out in the world) of the two specific inputs being evaluated (Bermúdez, 2003; Bohn, Call, & Völter, 2020; Völter & Call, 2017). Bermúdez (2003) was the first to argue that contrariety is a phylogenetic precursor of negation, terming it 'protonegation'. He argued that animals that cannot represent NEGATION might nevertheless learn through experience that specific pairs of predicates or states of the world are incompatible—EMPTY-FULL, PRESENT-ABSENT, VISIBLE-INVISIBLE, and so on—but might still never form a domain-general concept that these pairs are CONTRARY. Following this idea, Bohn, et al. (2020) state that what differentiates contrariety from negation is whether the operation is domain-general, leading them to propose a diagnostic for the possession NEGATION—if an animal can compare arbitrary inputs with which it has had no prior experience to determine that they are contraries, then on their view the animal has the concept NEGATION. Reviewing studies of non-human animals, Bohn, et al. (2020) argue that, although many studies have found evidence of individuals from many species reasoning about specific contrary pairs in specific contexts, no studies have provided conclusive evidence of any non-human animal having the capacity to represent contrariety as a domain-general operator, in the absence of prior experience with the specific contraries. They argue, therefore, that there is not clear evidence that any non-human animal thinks negated thoughts.

We think that this view of the relation between negation and contrariety conflates the two different ways of being a precursor: *partial function* and *domain specificity*. As we argued above, there is no reason to assume that an operator with the partial function of negation must also be domain specific. The representation of contrariety is not necessarily any less abstract or domain-general than negation, since both negation and contrariety are logical functions that can in principle operate over diverse content, including representations of any states of the world. The distinction between an operator satisfying only the LNC, as contrariety does, and both LNC and LEM, as negation does, is orthogonal to whether or how much the inputs of the operators are restricted, and therefore also orthogonal to whether any animal (human or not) has prior

experience with the incompatibility of the inputs being compared. If partial function and domain specificity are separable, and contrariety is defined specifically as having partial function, then prior experience with the incompatibility of particular states cannot be necessary to represent contrariety, nor can the ability to compare contraries in the absence of prior experience distinguish contrariety from negation.

Moreover, it is unclear how any animal could learn that a pair of predicates are contraries if it could not represent them *as* contrary—that is, if it did not have access to the concept CONTRARY, which could take two predicates as its arguments and return true if they are contrary and false otherwise. Bermúdez (2003) assumes that representations of contrary pairs come from experience, but how could an animal interpret that experience in terms of that relation between two predicates, if it could not already represent the contrariety relation? Not only is it not necessary to have prior experience with the incompatibility of being in New Jersey and being in outer space to represent them as contrary, it is not even clear what could constitute relevant prior experience with the incompatibility of these states. What could count as evidence that one cannot be in both of those places at once? Certainly, just being in one of the two places cannot provide any information about the (im)possibility of being in the other. Evidence of contrariety is no more inherent in the world than evidence of negation, since no percept carries any information about what other percepts it is incompatible with, any more than any percept carries information about what it is not.

In short, contrariety is no less an example of a logical concept than negation; it is just a distinct logical function that plays some, but not all, of the functional role of negation. This, however, does not make CONTRARIETY any worse of a candidate to be a developmental precursor of the concept NEGATION. It is just the sense of *partial function* in which it might serve as a precursor. Critically, however, and contra Bermúdez (2003), Völter & Call (2017), and Bohn, et al. (2020), we do not see any way to build a fully domain-general CONTRARIETY concept out of experience with specific contrary pairs without presupposing that this concept is already deployed in representing the relation within each pair. To serve as precursor to NEGATION, it seems to us that domain-general CONTRARIETY must itself already be available.

5.2 How to build negation out of contrariety

For contrariety to be a developmental precursor of negation, there must be some way to start with a function that obeys only LNC and end up, through some developmental process, with a function that obeys both LNC and LEM. How could that work? Consider that, as Geach (1972) observed, contrariety is not a proposition-generating function in the same way that negation is, because it does not identify a unique proposition. While there is only one negation of a proposition p, there may be many different contraries (*you are in outer space,* or *in New Jersey,* or *in France,* or...). If p is true, every one of its contraries must be false; conversely, if any one of the contraries of p is true, then p must be false. Thus, the 'middle' that is excluded by the negation of any proposition p is the empty set at the intersection of p and the set of all propositions q such that q is a contrary of p. This gives us a definition of negation:

1. $\sim p := \{q|\, \text{contrary}(q, p)\}$

This can be read as defining the negation of p as the set of all q such that q is contrary to p. This way of formulating the relation between contrariety and negation suggests a developmental pathway from the former to the latter: negation could develop as a generalization from contrariety computed over particular arguments (all of which must be represented *as* instances of contrariety) to the set of all contraries. That, in turn, casts a new light on Bohn and colleagues' (2020) proposal of domain-general contrariety as the criterion for the possession of the concept NEGATION. We argued above that this criterion is not sufficient because, while an animal may be able to represent a domain-general concept CONTRARY that allows comparing any two arguments, that representation will still only have the partial function of NEGATION. However, the above definition of negation suggests that, though not sufficient, representations of domain-general contrariety are necessary to represent negation. In this light, Bohn and colleagues' (2020) conclusion that no non-human animal has been shown to represent NEGATION because none has been shown to entertain domain-general contrariety is valid.

5.3 Do infants think negated thoughts?

Returning to ontogenesis, is there evidence that infants can represent domain-general contrariety prior to their understanding that *no* and *not* mean NEGATION?

One test of infants' ability to think negated thoughts is their ability to reason by exclusion: given multiple options about the state of the world and subsequent evidence that one of those options is false, can infants exclude it from consideration to pursue the remaining

option(s)? This question has received a large amount of attention in recent investigations, with conflicting results. On the one hand, Cesana-Arlotti, et al. (2018; 2020) argue that infants as young as 12 months are able to represent negation, or at least a precursor of negation. Though they do not discuss what this precursor could be like, contrariety would fit the bill. In their tasks, infants watch a movie that starts with two objects: e.g., a snake and a ball. In the crucial Inference condition, a screen rises to cover both objects. Next, a cup appears, swoops in behind the screen, and comes out with the top of an object peeking out over its rim. The snake and ball are designed to look identical when only their tops are visible, so a viewer cannot tell which object is in the cup, and so might set up a representation of a disjunction, *either the ball or the snake is in the cup*. Next, a snake briefly emerges from the side of the occluder before retreating behind it again. Representing that the snake is behind the occluder could be taken as contrary to the snake being in the cup, licensing the elimination of that option from the disjunction, and the inference that it is the ball that is in the cup. In a control No-Inference condition, there was no ambiguity about the locations of either the ball or the snake throughout the whole event.

Cesana-Arlotti et al. used several sources of evidence to argue that infants reasoned by exclusion. First, they found that infants looked longer at the final inconsistent events than the final consistent events in both conditions: either a ball emerging from behind the occluder (implying three objects in total), or a snake being revealed in the cup (implying two snakes and no ball). However, as Cesana-Arlotti and colleagues acknowledge, these results do not necessarily mean infants were reasoning by exclusion. Infants' longer looking could reflect their attempting to reconcile the initial representation of a snake and a ball, held in working memory, with the final evidence either for two snakes or for three total objects.

Stronger evidence for reasoning by exclusion comes from comparing patterns of eye movements and pupil dilation between the Inference and No-Inference conditions during the period in which an inference might have been made. Just after the snake emerged from behind the occluder, infants' pupils dilated more in the Inference condition, suggesting greater attention or processing, and they shifted their gaze more from the occluder to the cup, consistent with their inferring its contents. Furthermore, more pupil dilation and shifts to the cup in the inference condition predicted greater looking times on the final violation of expectancy test trials. Cesana-Arlotti, et al. (2018) take these data to show both that infants eliminated the possibility that the snake was in the cup (the function of negation) and that the infants concluded that it is the ball

that must therefore be in the cup (the function of disjunction). They found these signatures in 12- and 19-month olds, and 14-month-olds also succeeded in a variant that further requires integrating the representation that results from this inference with a representation of agents' preferences (Cesana-Arlotti et al., 2020).

These studies provide strong evidence that infants are inferring the identity of the object in the cup in the inference condition. However, Feiman, Mody, & Carey (2022) have challenged the conclusion that this inference depends on infants deploying the concepts NEGATION or DISJUNCTION. They argue that there is an alternative explanation of these results that involves infants' ability to maintain a 1-to-1 mapping between the spatiotemporally distinct individuals and the kinds of objects they have seen.

On this alternative, infants first create a mental model of the initial scene and hold it in working memory: e.g., one snake and one ball on the stage. As the scene unfolds, they simply build a perceptual representation of what they are looking at, monitoring for consistency between the unfolding scene and the initial model. When the objects are occluded and the cup scoops one up, the perceptual representation of what is in the cup now contains a 'bare' object file (a spatiotemporal index with no other properties; see Scholl, 2001), or perhaps an object with a red top part. This scene is still consistent with the initial model. At this point, infants need not represent any specific alternatives about what is in the cup or what is behind the screen. Next, they see the snake emerge from behind the screen, and they specify the current location of the snake that had been held in their working memory model. Now infants are in a position to make an inference. Finding out where the snake is leaves only two unknowns that just happen to match each other—an object with an unknown location (the ball that was present at the start, a representation of which is still held in working memory), and a location (the cup) visibly containing an indexed object, which also happens to look like that ball. This allows for a 1-to-1 mapping of the spatiotemporal individual to the kind—*that object in the cup is the ball*—producing a genuine inference, which may result in corresponding increases both in looks to the cup and in pupil dilation. Critically, however, on this account infants never wonder whether what is in the cup is the ball *or* the snake, and they never consider the possibility that the snake is in the cup (let alone exclude this possibility from consideration once they see the snake emerge from behind the screen). That is, they represent neither negation nor disjunction. The appearance of the snake is only incidentally informative about the location of the ball because it happens to

leave only one way to specify the two remaining unknowns (where the ball is, and what is in the cup; see also Jasbi, et al., 2019).

Consistent with this alternative, Feiman, et al. (2022) find that infants fail on two other tasks of reasoning by exclusion until they are 17 months old—5 months later than Cesana-Arlotti, et al.'s findings of success, but still more than 6 months before they have mapped *no* and *not* to NEGATION. In one task, infants are shown a toy being hidden behind an occluder in one of two buckets. The experimenter then removes the occluder, picks up one of the buckets, and turns it upside down facing the infant to show that it is empty, before finally placing it back where it was. Infants are invited to search for the hidden toy. In a second task using a 'blicket detector' design (Gopnik, et al. 2004), infants are shown two differently colored blocks being placed on a box, causing a spinning light inside the box to 'magically' activate. They are then shown that one of these two blocks alone does not activate the magic box. Finally, they are invited to choose between the two blocks—which should they try to place on the box to make it light up? Feiman, et al. (2022) found the same developmental trajectory in both tasks: 15-month-olds failed, choosing between the two options at chance as if they had never witnessed the negative evidence. 17-month-olds succeeded, though their performance was not significantly better than the younger infants', and 19-month-olds succeeded more robustly, significantly outperforming 15-month-olds.

This evidence suggests that, at least by 19 months, infants robustly reason by exclusion using particular contraries: EMPTY-FULL in the search task, and ACTIVE-INERT in the blicket detector task. Moreover, the convergence in ages of success and failure between the two tasks suggests that the ability to represent contrariety could be domain-general, independent of the particular predicates being compared. Still, co-incidences in ages of success between tasks could just be coincidences; after all, there are many tasks on which 19-month-olds outperform 15-month-olds, and surely not all of them rely on the development of just one ability. More direct tests of 19-month-olds' capacities await. Future research may look for additional evidence of within-individual correlations between different tests, asking whether children who succeed on one task are more likely to succeed on the others, controlling for shared performance factors like attention and inhibitory control. Even better evidence would be if success or training on one task causally affects performance on another—a kind of structural priming approach (Pickering & Ferreira, 2008) applied to the study of infants.

While there is yet no more converging evidence to suggest that 19-month-olds represent contrariety domain-generally, there is additional evidence that younger infants do not. Feiman, Carey, and Cushman (2015) asked whether infants would represent a generalization that an agent prefers *something else* over a particular object. Seven- and 14-month-old infants were shown a hand choosing between two objects: A (e.g, a ball) versus a sequence of other alternatives B, C, D, etc. (e.g., a brush, a bottle, a stick, etc.) In a condition in which the hand always reached for A, even 7-month-olds expected it to continue reaching for A, (the ball) over novel alternatives. But when either 7- or 14-month-old infants were shown the hand always reaching for different novel objects rather than the ball, they did not expect the hand to reach for the next novel object when paired with the ball. This is consistent with an inability to represent or condition an expectation on a contrary representation involving a variable—*contrary(x, ball)*—i.e., a failure to represent that the hand will reach for *something contrary to the ball* on subsequent trials.

Although not done with this purpose in mind, Benavides-Varela & Mehler (2015) also explored whether infants can condition a response on a representation of *something contrary to A*. They found that 7-month-olds learned that one novel word (A) predicted something would happen on the right side of the screen, but failed to learn that any of 12 other words (which might be represented as either 12 different specific contraries to A, or simply as 12 different instances of *something contrary to A*) predicted left. Future research might explore whether infants can represent a variable *contrary(x, A)* at the same age as they succeed on the search and blicket detector tasks of Feiman, et al. (2022), or whether the ability to consider an arbitrary variable as one of the contraries emerges later in development, and might require an additional piece of computational machinery.

Adding a further degree of abstraction, Hochmann, Mody and Carey (2016) provide positive evidence that infants under 15 months of age can formulate a rule involving a match computation over variable inputs, *match(x,x)*, but not one involving arbitrary incompatible contraries, *contrary(x,y)*. Hochmann et al. trained one group of 14-month-old infants on a visual Match-To-Sample (MTS) task and a different group on a Non-Match-To-Sample (NMTS) task, and showed that infants succeeded at both. In both tasks, participants are shown three cards with novel shapes on each trial: one with a sample shape in the middle, and then two with other shapes on the sides, one of which is identical to the sample. Shortly thereafter, in MTS, the shape that matches the sample begins to spin, accompanied by a sound. In NMTS, the shape that

mismatches the sample does so. In both conditions, infants learned to anticipate which shape would spin after 12 trials, looking at the correct shape before it moved. Success in NMTS seemingly provides evidence for conditioning a response on a contrary (choose the shape that mismatches the sample). However, a subsequent test ruled out this conclusion. After training on either MTS or NMTS, infants were again shown the sample and two choice cards on either side, but both choice cards were flipped over, such that the baby could not see what shapes were on them. Then the shape on only one of the choice cards was briefly made visible; and then rehidden. The shape on the other card was never revealed. On half of the trials the choice that had been revealed matched the sample; on the other half, it mismatched. When they had seen the *matching* stimulus, participants in both the MTS and NMTS condition chose the correct location, the card with the same shape in MTS and the card with the (never seen) mismatching shape in NMTS. That is, they could follow both of the rules *seek match* (in MTS) and *avoid match* (in NMTS). However, when only the mismatching shape had been revealed, infants looked at chance between the two cards, both in MTS and NMTS. Fourteen-month-olds had not formulated the rule *seek contrary* (in NMTS) or the rule *avoid contrary* (in MTS). In other words, *contrary* was not part of the computation underlying success in NMTS; only *match* was. The same results have been found in pigeons (Zentall, Andrews, & Case, 2018), who, like infants 15 months of age and younger, also fail tests of reasoning by exclusion (Aust, et al., 2008, Cumming & Berryman, 1961; but see Lauffer, et al., 2017).

Relatedly, Hochmann, Carey, & Mehler (2018) also attempted to teach 7- and 12-month-old infants two rules: that two stimuli that were the same as each other predicted that an object would appear to the right, while two stimuli that were not the same as each other predicted an object appearing to the left. Again, the specific stimuli changed on each trial. And again, at both ages, infants learned the generalization *match (x, x) predicts right*, but performed at chance with the generalization *contrary (x, y) predicts left*.

Both of these tasks require a more abstract representation of contrariety because the specific stimuli that the infants see change on every trial. To succeed, infants must formulate a rule over pairs of shapes they have seen and apply that rule to novel stimuli they had never compared for contrariety before. The failure of infants under 15 months of age at these tasks is consistent with the failure of infants under 17 months of age on the two tasks reported by Feiman, et al. (2022), in which infants only had to represent contrariety over specific, fixed

inputs. It is also consistent with the alternative explanation of 12-month-old infants' success on Cesana-Arlotti, et al.'s (2018; 2020) tasks, suggesting that infants are not relying on contrariety (let alone negation) in those tasks. The question of whether infants as young as 17- to 19-months of age would succeed on tasks that require representing contrariety with one or two variable inputs—*contrary(x, A)* as in Feiman, et al. (2015) and Benavides-Varela, et al. (2015), or *contrary(x, y),* as in Hochmann, et al. (2016)—is presently open.

In sum, a growing body of evidence suggests that both non-human animals (Bohn, et al., 2020) and human infants at least by 17 to 19 months of ages (Feiman, et al., 2022) if not already by 12 months (c.f. Cesana-Arlotti, et al., 2018; 2020) can represent specific states of the world as contrary and reason on that basis, such as the ball being present in a bucket and the ball being absent from that bucket. There is currently only weak evidence—coincidence in the ages of success between the two tasks reported by Feiman, et al. (2022)—that infants' representations of contrariety at that age are domain-general. The open questions are, first, what concepts do infants between 17 and 19 months have that younger infants lack? Is it a domain-specific or domain-general concept CONTRARY, or is it perhaps full-fledged NEGATION (a representation of the set of all contraries)? Second, how do infants' capacities change between the ages at which they fail on tasks that require (at minimum) representations of contrariety, prior to 17 months (Benavides-Varela & Mehler, 2015; Feiman, et al., 2022; Feiman, et al., 2015; Hochmann, et al., 2016; 2018) and the age at which they first succeed on such tasks? If they employ either a domain-general or a more limited concept CONTRARY on these tasks, how does that emerge? While we have discussed how negation could develop from contrariety, all we have said about the potential emergence of contrariety itself is that accounts of its emergence from specific experience with contrary pairs seem question-begging. We take this to be a major unsolved question. Third, if 17-month-olds do not yet have NEGATION, what (if anything) changes in infants' conceptual repertoire from their first unambiguous successes on tasks involving contraries at 17 months and their mapping of NEGATION to its expression in their first language by about 24 to 27 months? Section 5.2 suggested how it might be possible to build negation out of the partial function of contrariety and Section 4 discussed how cross-linguistic differences can be leveraged to test whether children build negation out of domain-specific precursors, such as rejection and nonexistence, but there is currently no evidence about what, if any, construction actually takes place.

6 Conclusion

Learning to express negation requires two distinct abilities: having the concept of negation to think with, and mapping this concept to the linguistic forms that express it in one's language. In Section 2, we argued that children's early uses of negative expressions reflect their having mapped more limited concepts to those expressions. In Section 3, we argued that the difficulty of the mapping problem is sufficient to account for the long delay between these first uses of *no*, which appear not to express negation yet, and their subsequent mapping of both *no* and *not* to the logical concept, which takes place between 6 and 12 months later. Nevertheless, there must be some conceptual development of NEGATION—if not in infancy, then over the course of evolution. In Section 4, we suggested how a certain kind of conceptual development in infancy—from more domain-specific precursors like REJECTION—might produce testable cross-linguistic differences in language acquisition. In Section 5, we distinguished domain-specific precursors from precursors that could be domain-general but have only part of the function of full-fledged negation, and considered how negation might be constructed from one domain-general precursor with a partial function—contrariety.

The resulting landscape leaves us with many open questions: what are the actual precursors of negation in thought? What are the origins of those precursors? Are they precursors in ontogenesis or phylogenesis? If in phylogenesis (that is, if the concept NEGATION is innate in humans), what limits infants' ability to use it in reasoning tasks until about 17 months of age? What, if anything, changes in infants' representation of negation after 17 months? Whether there is conceptual development or not, how exactly do toddlers then solve the complex mapping problem of connecting the concept of negation with its expressions in their language? These questions seem to us hard, but also specific and tractable directions for future work.

7 References

Aust, U., Range, F., Steurer, M., & Huber, L. (2008). Inferential reasoning by exclusion in pigeons, dogs, and humans. *Animal Cognition*, *11*(4), 587–597. https://doi.org/10.1007/s10071-008-0149-0

Austin, K., Theakston, A., Lieven, E., & Tomasello, M. (2014). Young children's understanding of denial. *Developmental Psychology*, *50*(8), 2061–2070. https://doi.org/10.1037/a0037179

Benavides-Varela, S., & Mehler, J. (2015). Verbal Positional Memory in 7-Month-Olds. *Child Development*, *86*(1), 209–223. https://doi.org/10.1111/cdev.12291

Bermúdez, J. L. (2003). *Thinking without words*. Oxford University Press.

Bloom, L. (1970). Syntactic and Semantic Development of Early Sentence Negation. In *Language Development: Form and Function in Early Grammars*. MIT Press.

Bohn, M., Call, J., & Völter, C. J. (2020). Evolutionary Precursors of Negation in Non-Human Reasoning. In V. Déprez & M. T. Espinal (Eds.), *The Oxford Handbook of Negation* (pp. 576–588). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780198830528.013.40

Cameron-Faulkner, T., Lieven, E., & Theakston, A. (2007). What part of *no* do children not understand? A usage-based account of multiword negation. *Journal of Child Language*, *34*(2), 251–282. https://doi.org/10.1017/S0305000906007884

Casenhiser, D. M. (2005). Children's resistance to homonymy: An experimental study of pseudohomonyms. *Journal of Child Language*, *32*(2), 319–343. https://doi.org/10.1017/S0305000904006749

Cesana-Arlotti, N., Kovács, Á. M., & Téglás, E. (2020). Infants recruit logic to learn about the social world. *Nature Communications*, *11*(1), 5999. https://doi.org/10.1038/s41467-020-19734-5

Cesana-Arlotti, N., Martín, A., Téglás, E., Vorobyova, L., Cetnarski, R., & Bonatti, L. L. (2018). Precursors of logical reasoning in preverbal human infants. *Science*, *359*(6381), 1263–1266. https://doi.org/10.1126/science.aao3539

Choi, S. (1988). The semantic development of negation: A cross-linguistic longitudinal study. *Journal of Child Language*, *15*(3), 517–531. https://doi.org/10.1017/S030500090001254X

Cumming, W. W., & Berryman, R. (1961). SOME DATA ON MATCHING BEHAVIOR IN THE PIGEON [1]. *Journal of the Experimental Analysis of Behavior*, *4*(3), 281–284. https://doi.org/10.1901/jeab.1961.4-281

Dautriche, I., Fibla, L., Fievet, A.-C., & Christophe, A. (2018). Learning homophones in context: Easy cases are favored in the lexicon of natural languages. *Cognitive Psychology*, *104*, 83–105. https://doi.org/10.1016/j.cogpsych.2018.04.001

de Carvalho, A., Crimon, C., Barrault, A., Trueswell, J., & Christophe, A. (2021). "Look! It is not a bamoule!" 18- and 24-month-olds can use negative sentences to constrain their interpretation of novel word meanings. *Developmental Science*. https://doi.org/10.1111/desc.13085

De Villiers, P., & De Villiers, J. (1979). Form and Function in the Development of sentence negation. *Papers and Reports on Child Language Development*, *17*(Aug).

Dimroth, C. (2010). The Acquisition of Negation. In *The Expression of Negation* (pp. 39–73). DE GRUYTER MOUTON. https://doi.org/10.1515/9783110219302.39

Drozd, K. F. (1995). Child English pre-sentential negation as metalinguistic exclamatory sentence negation. *Journal of Child Language*, *22*(3), 583–610. https://doi.org/10.1017/S030500090000996X

Feiman, R., Carey, S., & Cushman, F. (2015). Infants' representations of others' goals: Representing approach over avoidance. *Cognition*, *136*, 204–214. https://doi.org/10.1016/j.cognition.2014.10.007

Feiman, R., Mody, S., & Carey, S. (2022). The development of reasoning by exclusion in infancy. *Cognitive Psychology*, *135*, 101473. https://doi.org/10.1016/j.cogpsych.2022.101473

Feiman, R., Mody, S., Sanborn, S., & Carey, S. (2017). What Do You Mean, No? Toddlers' Comprehension of Logical "No" and "Not." *Language Learning and Development*, *13*(4), 430–450. https://doi.org/10.1080/15475441.2017.1317253

Feiman, R., Zimmerman, S., & Snedeker, J. (2019). *The acquisition of negation: International adoption pulls apart conceptual and linguistic limits on learning*. The 44th Annual Boston University Conference on Language Development.

Feiman, R., Zimmerman, S., & Snedeker, J. (in prep). *International adoption separates cognitive development and language knowledge in the acquisition of negation*.

Fenson, L., Marchman, V. A., Thal, D. J., Dale, P. S., Reznik, J. S., & Bates, E. (2007). *MacArthur-Bates Communicative Development Inventories: User's Guide and Technical Manual (2nd ed.)* (2nd ed.). Brookes Publishing Co.

Floyd, S., & Goldberg, A. E. (2021). Children make use of relationships across meanings in word learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *47*(1), 29–44. https://doi.org/10.1037/xlm0000821

Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2017). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, *44*(3), 677–694. https://doi.org/10.1017/S0305000916000209

Geach, P. T. (1972). *Logic matters*. University of California Press.

Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, *73*, 135–176.

Gleitman, L. R., Cassidy, K., Nappa, R., Papafragou, A., & Trueswell, J. C. (2005). Hard Words. *Language Learning and Development*, *1*(1), 23–64. https://doi.org/10.1207/s15473341lld0101_4

Gleitman, L. (1990). The Structural Sources of Verb Meanings. *Language Acquisition*, *1*(1), 53.

Gomes, V., Doherty, R., Smits, D., Goldin-Meadow, S., Trueswell, J., & Feiman, R. (2021). *Not as hard as it looks? Linguistic, and not conceptual, ability rives negation acquisition*. The 46th Annual Boston University Conference on Language Development.

Gomes, V., Doherty, R., Smits, D., Goldin-Meadow, S., Trueswell, J., & Feiman, R. (in prep). *Not as hard as it looks? Linguistic, and not conceptual, ability rives negation acquisition*.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A Theory of Causal Learning in Children: Causal Maps and Bayes Nets. *Psychological Review*, *111*(1), 3–32. https://doi.org/10.1037/0033-295X.111.1.3

Gopnik, A., & Meltzoff, A. N. (1985). From people, to plans, to objects. *Journal of Pragmatics*, *9*(4), 495–512. https://doi.org/10.1016/0378-2166(85)90018-9

Hochmann, J.-R., Carey, S., & Mehler, J. (2018). Infants learn a rule predicated on the relation same but fail to simultaneously learn a rule predicated on the relation different. *Cognition*, *177*, 49–57. https://doi.org/10.1016/j.cognition.2018.04.005

Hochmann, J.-R., Mody, S., & Carey, S. (2016). Infants' representations of same and different in match- and non-match-to-sample. *Cognitive Psychology*, *86*, 87–111. https://doi.org/10.1016/j.cogpsych.2016.01.005

Horn, L. R. (1989). *A natural history of negation*. University of Chicago Press.

Hummer, P., Wimmer, H., & Antes, G. (1993). On the origins of denial negation. *Journal of Child Language*, *20*(3), 607–618. https://doi.org/10.1017/S0305000900008503

Jasbi, M., Bohn, M., Long, B., Fourtassi, A., Barner, D., & Frank, M. C. (2019). Comment on Cesana-Arlotti et al. (2018). https://doi.org/10.31234/osf.io/g2h7m

Lauffer, M. C., Castro, L., & Wasserman, E. A. (2017). Chrysippus's pigeon: Exclusion-based responding in an avian model. *Journal of Experimental Psychology: Animal Learning and Cognition*, *43*(2), 139–146. https://doi.org/10.1037/xan0000131

Macnamara, J. (1986). *A Border Dispute: The Place of Logic in Psychology*. The MIT Press.

Mazzocco, M. M. M. (1997). Children's interpretations of homonyms: A developmental study. *Journal of Child Language*, *24*(2), 441–467. https://doi.org/10.1017/S0305000997003103

Nordmeyer, A. E., & Frank, M. C. (2018). Individual variation in children's early production of negation. *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, 6.

Pea, R. D. (1980a). LOGIC IN EARLY CHILD LANGUAGE*. *Annals of the New York Academy of Sciences*, *345*(1), 27–43.

Pea, R. D. (1980b). The Development of Negation in Early Child Language. In *The Social foundations of language and thought: Essays in honor of Jerome S. Bruner* (pp. 156–186). W. W. Norton Company.

Peters, A., & Zaidel, E. (1980). The acquisition of homonymy. *Cognition*, *8*(2), 187–207. https://doi.org/10.1016/0010-0277(80)90012-8

Pickering, M. J., & Ferreira, V. S. (2008). Structural priming: A critical review. *Psychological Bulletin*, *134*(3), 427–459. https://doi.org/10.1037/0033-2909.134.3.427

Schachter, P., & Otanes, F. T. (1983). *Tagalog reference grammar*. University of California Press.

Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, *80*(1–2), 1–46. https://doi.org/10.1016/S0010-0277(00)00152-9

Snedeker, J., Geren, J., & Shafto, C. L. (2007). Starting Over: International Adoption as a Natural Experiment in Language Development. *Psychological Science*, *18*(1), 79–87. https://doi.org/10.1111/j.1467-9280.2007.01852.x

Snedeker, J., Geren, J., & Shafto, C. L. (2012). Disentangling the effects of cognitive development and linguistic expertise: A longitudinal study of the acquisition of English in internationally-adopted children. *Cognitive Psychology*, *65*(1), 39–76. https://doi.org/10.1016/j.cogpsych.2012.01.004

Veselinova, L. (2013). Negative existentials: A cross-linguistic study. *Rivista Di Linguistica*, *25*(1), 107–145.

Völter, C. J., & Call, J. (2017). Causal and inferential reasoning in animals. In J. Call, G. M. Burghardt, I. M. Pepperberg, C. T. Snowdon, & T. Zentall (Eds.), *APA handbook of comparative*

psychology: Perception, learning, and cognition. (pp. 643–671). American Psychological

   Association. https://doi.org/10.1037/0000012-029

Zentall, T. R., Andrews, D. M., & Case, J. P. (2018). Sameness May Be a Natural Concept That Does

   Not Require Learning. *Psychological Science*, *29*(7), 1185–1189.

   https://doi.org/10.1177/0956797618758669