Gomes, V., Doherty, R., Smits, D., Goldin-Meadow, S., Trueswell, J. C., & Feiman, R. (2023). It's not just what we don't know: The mapping problem in the acquisition of negation. *Cognitive Psychology, 145*, 101592.

It's not just what we don't know: The mapping problem in the acquisition of negation

Victor Gomes[a*], Rebecca Doherty[bd], Daniel Smits[b], Susan Goldin-Meadow[c], John C. Trueswell[a], Roman Feiman[b]

[a]Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA
[b]Department of Cognitive, Linguistic & Psychological Sciences, Brown University, Providence, RI, USA
[c]Department of Psychology, University of Chicago, Chicago, IL, USA
[d]Department of Psychology, University of Bath, UK

*Corresponding author at: Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA.
E-mail address: vgomes@upenn.edu (V. Gomes).

Abstract

How do learners learn what *no* and *not* mean when they are only presented with what is? Given its complexity, abstractness, and roles in logic, truth-functional negation might be a conceptual accomplishment. As a result, young children's gradual acquisition of negation words might be due to their undergoing a gradual conceptual change that is necessary to represent those words' logical meaning. However, it's also possible that linguistic expressions of negation take time to learn because of children's gradually increasing grasp of their language. To understand what *no* and *not* mean, children might first need to understand the rest of the sentences in which those words are used. We provide experimental evidence that conceptually equipped learners (adults) face the same acquisition challenges that children do when their access to linguistic information is restricted, which simulates how much language children understand at different points in acquisition. When watching a silenced video of naturalistic uses of negators by parents speaking to their children, adults could tell when the parent was prohibiting the child and struggled with inferring that negators were used to express logical negation. However, when provided with additional information about what else the parent said, guessing that the parent had expressed logical negation became easy for adults. Though our findings do not rule out that young learners also undergo conceptual change, they show that increasing understanding of language alone, with no accompanying conceptual change, can account for the gradual acquisition of negation words.

Keywords: Language Acquisition, Negation, Human Simulation Paradigm, Conceptual Development, Logical Reasoning, Word Learning

1. Introduction

"Contrariwise," continued Tweedledee, "if it was so, it might be; and if it were so, it would be; but as it isn't, it ain't. That's logic."

- Caroll, 1872

How do children learn the meanings of words like *no* and *not*? It is tempting to think that word learning in general is accomplished simply by pairing word occurrences with their situational contexts of use, relying on the world to constrain word meanings. But negation is nowhere to be found in the world (see Russell, 1954), and yet children still manage to learn *no* and *not* long before they ever take an introductory logic course. These words, which are commonly spoken by parents to their children, have a message-internal meaning that corresponds (roughly, see Horn, 1989) to negation – formally, a unary operator that inverts the truth-value of a statement. Because negation isn't inherent in the world, thinking with negation depends on how the world is construed – and there are always alternative construals (Gleitman & Gleitman, 1992). For example, aiming to hit a window with a ball but missing it and hitting the wall instead, could be construed either as *not hitting the window*, or as a corresponding affirmative, *hitting the wall*, which involves no negation (Ackrill, 1975).

Nevertheless, speakers do (at least sometimes) talk about the world around them. When that happens, the world can constrain learners' hypotheses about the construal that the speaker wants to convey, and hence about the meanings of the words used to convey it. If word learning requires reverse-engineering the speaker's construal, then words whose meanings depend on specific, atypical, or newsworthy construals of the world should be delayed in entering a learner's vocabulary. Consistent with this view, words whose meanings are less imageable tend to be learned later by children, independent of cognitive ability (Snedeker, Geren & Shafto, 2007) and are harder to learn when the only available information is the surrounding visual environment, with no accompanying language (Gillette, Gleitman, Gleitman & Lederer, 1999). The meanings of these 'hard' words are not obvious because they do not easily align with a learner's construal of the world when hearing them, making them difficult to learn from observation of the world alone. To learn them, learners might need additional message-internal clues (e.g., co-occurring words and co-occurring syntax whose meanings the learner has already acquired, as in syntactic bootstrapping, Gleitman, 1990; Gleitman, Cassidy, Nappa, Papafragou & Trueswell, 2005).

Here we explore how this view of word learning can help explain how children learn what might be the least imageable words – words that express negation (*no, not*) – from the available evidence, i.e., from the situational contexts and speech they encounter in the home. Below, we first discuss what is known about children's acquisition of these words, which suggests that children piece together their meanings only gradually. We then consider two mutually compatible, but alternative explanations for this developmental trajectory: either that young children lack the ability to construe the world through negation, or that they lack the information they need to infer that speakers who use *no* and *not* are construing the world that way. We then offer novel experimental results in support of the latter alternative being a sufficient explanation on its own.

1.1. What do children *no* and when do they *no* it?

In English, the two words that express negation are *no* and *not* (with its close morphological relative, *-n't*). Children begin producing *no* first (50% of parents report that their children say *no* by 15 months; Frank et al., 2016), and only begin producing *not* about a year later, around 27-28 months (Frank et al., 2016). Along with changes in the lexical forms they produce, Bloom (1970) also found changes in what children use *no* and *not* to express. Studying how three children between 19 and 24 months of age interacted with their caregivers, Bloom proposed a tripartite taxonomy of emerging uses of negation in children's speech: Nonexistence, Rejection, and Denial. Bloom found that children's primary use of negation early on (between 19-22 months of age) was to Reject offers or commands (e.g., "No veggies!", when being offered veggies) and to comment on the Nonexistence of something they expected to be there (e.g., "No more!", when noticing that they had drunk all the juice). Only later (at about 2.5 years of age) did children start to Deny the truth of propositions (e.g., "That's not a cat!"; Bloom, 1970). Denials appeared around the same time children started to say *not* in addition to *no*, and were expressed by both forms. Since Bloom, more recent studies of other children's speech have replicated this ordering (Choi, 1988; Nordmeyer & Frank, 2017a) and added other categories to the taxonomy, while also finding that Rejection and Denial are by far the most frequent functions in children's speech.[1]

Do these changes in what children say reflect changes in their understanding of what *no* and *not* mean? Bloom (1970) noted that it may be due to Denial requiring representing the corresponding affirmative, and so holding two propositions in mind at once (see also McNeill & McNeill, 1967). However, Pea (1980) first raised the possibility that, although Denial uses of *no* and *not* are clear expressions of truth-functional negation, Rejection and Nonexistence uses might reflect different, more limited concepts. More recent evidence suggests that Pea may have been right; children's comprehension of Denials matches their production, and emerges well after they first

[1] Choi (1988) expanded this taxonomy with additional functions of negation, including Prohibition. As Choi (1988) notes, Rejection and Prohibition largely differ based on the hierarchy an interaction takes place within (children can't prohibit their parents from doing something), with parents uttering Prohibitions much more than Rejections for this reason. We will use a merged Bloom-Choi taxonomy of negation for our analyses of parental utterances: Denial, Nonexistence, and Prohibition (Bloom, 1970; Choi, 1988).

start producing *no* to express Rejection or Nonexistence (Austin, Theakston, Lieven & Tomasello, 2014; Feiman, Mody, Sanborn & Carey., 2017; though, for earlier evidence of Denial comprehension, see de Carvalho, Crimon, Barrault, Trueswell, & Christophe, 2021). For example, Austin et al. (2014) and Feiman et al. (2017) found that children younger than 2-years-old did not understand when either *no* or *not* were used to express logical negation in a search task. In these studies, children played a game in which they had to find a ball that was hidden in one of two containers (e.g., a bucket or a truck). On some trials, the children were told where the ball was, and on other trials, they were told where it was not. In one condition, the negative information was provided with the word *no* (e.g., the parent asked, “Is it in the bucket?” and the experimenter who had hidden the ball said, “no” (Austin, et al., 2014) or “no, it’s not” (Feiman, et al., 2017). In another condition, the experimenter used *not* exclusively (e.g., “It’s not in the bucket”, in both studies). Understanding *no* or *not* should lead children to search in the other location, the truck. Both studies found, however, that children only understood the negative information (*no* and *not*) around 2.5 years of age, around the age when children typically begin to produce Denials. Both studies also found that children began to understand this use of both *no* and *not* at the same age (despite saying *no* at much younger ages), suggesting a common limit on the acquisition of the Denial meaning for both words. Feiman et al. (2017) and Feiman, Mody, & Carey (2022) also presented children with information of absence visually (showing them the empty bucket), and found them successfully searching in the other container a full year earlier, suggesting that their difficulty was specific to understanding the words *no* and *not*, rather than to reasoning by exclusion in this task.

These findings suggest that Denial uses of *no* express a distinct meaning from Rejection uses. Children under the age of two produce *no* regularly to express Rejection, at the same time that they succeed on non-verbal tasks that involve reasoning from an object’s absence to its location elsewhere (Feiman, Mody, Carey, 2022). Still, despite their ability to reject, by saying *no,* they fail to understand what it means to answer “no” in response to the question, “Is the ball in the bucket?” As Feiman, et al. (2017) argued, this suggests that Rejection not only precedes Denial in children’s production, but in comprehension as well. Moreover, in some languages, such as Korean (Choi, 1988) and Tagalog (Schachter & Otanes, 1983), Rejection and Denial are expressed by morphologically unrelated forms. To the extent that different words tend to express different meanings, this suggests that these meanings are distinct for adults as well as for children. Thus, children’s changing use of *no* likely captures changes in the meanings they have acquired for both *no* and *not,* and not merely differences in their use of a single meaning.

While there is significant evidence that Rejection and Denial express different meanings, it is less clear whether Nonexistence uses of *no* and *not* reflect a separate meaning from truth-functional negation or whether they are more like a very limited use-case of Denial. Consistent with the latter possibility, Bloom (1968) originally found that Nonexistence was less consistently produced early on than the other functions she identified, while in a case study, Cameron-Faulkner, et al. (2007) found it was rare (~5% of negations, across ages) in one child's input. Later corpus analyses similarly found Nonexistence uses to be vanishingly rare in child productions (Nordmeyer & Frank, 2018a) and, in both of the corpora we use for the present study, Nonexistence was far less common in parents' speech than Prohibition and Denial. On the other hand, although the usage of *no* and *not* to describe the ball not being in the bucket (Austin, et al., 2014; Feiman, et al., 2017), could be interpreted as expressing Nonexistence as well as truth-functional negation, both of these studies found children failing to find the ball 10 months after Bloom (1968) reports Nonexistence first emerges in production. This suggests that *no* and *not* in this task do not, in fact, get interpreted as expressing the same meaning that children might already be producing. Unlike children's Nonexistence productions, these tasks may require having mapped *no* and *not* to truth-functional negation (see Feiman, et la. 2017 for discussion). Additionally, Veselinova (2013) found that 42 languages, approximately half of those she studied, express Nonexistence using entirely distinct lexical forms (i.e. different negators) from truth-functional negation. Whether Nonexistence expresses truth-functional negation or a different meaning, its relative rarity in both children's and parents' speech suggests that its acquisition may not be a crucial milestone or ingredient in the acquisition of truth-functional negation.

Given that Denial uses are the clearest expression of truth-functional negation, why is learning to use *no* and *not* to express Denial challenging for the youngest learners? Is it because they cannot yet think logically negated thoughts as early as they can think thoughts that involve refusing or rejecting, and are therefore unable to hypothesize negation as the meaning of a word (even of *no*, which they already say)? Or is it simply because, like a stranger in a strange land, they haven't yet learned that *no* and *not* mean negation? Perhaps expressions of logical negation are 'hard words' in a way that expressions of rejection are not. Children cannot look out into the world and see the ball "not hitting the window," because that is indistinguishable from it "hitting the wall." But maybe they can look out and see someone expressing aversion or a desire for some event to stop. However, with additional linguistic information ("It did ___ hit the window"), maybe learners can also realize that someone means to be conveying a negated thought with that missing word, conceptualizing it as a negator. Following Gillette et al. (1999), we will call the former hypothesis the Conceptual Change account, and the latter the Information Change account. Below we unpack these two accounts further and test between them empirically.

1.2. Conceptual Change: The case of the missing concept

One explanation for the delay in the production and comprehension of truth-functional negation is that it is a result of a limited conceptual repertoire or still-developing cognitive ability. The problem, then, is not simply that a young learner does not know what *not* means in their language, but that they could not conceive of its meaning because they cannot think any thoughts that require truth-functional negation. In this case, younger children do not map truth-functional negation to *no* or *not,* and do not produce Denials, because they either do not have, or are incapable of engaging, the requisite concept.

Since Pea (1980), conceptual explanations like these have been the dominant account. As Pea states it: "Negation is a semantic domain which readily lends itself to conceptual analyses of representational complexity. The claim is that for the child to conceive the different meanings of negation…increasingly abstract forms of cognitive representation are required." Under this view, Rejection and Nonexistence precede Denial because of their relative conceptual simplicity. For example, as Dimroth (2010) writes, reviewing the literature, "Rejection typically applies to holistic actions or events and the topic of this type of negation has no need for internal representation, because it is always present in the context." On the other hand, Denial as an expression of truth-functional negation would be acquired later because of its much more complex conceptual requirements; children must "simultaneously represent two mental models," one of the true state of affairs and the other of the false counterpart (Dimroth, 2010; Pea, 1980; Hummer et al., 1993).

One challenge for this view is mounting evidence that infants may already have the ability to think negated thoughts before they start either producing or understanding verbal expressions of negation. Eleven-month-old infants can learn to expect sequences of syllables in which the last syllable is *not* the same as the preceding (AAA¬A, where the specific identities of A and ¬A vary across trials; Hochmann & Toro, 2021). Infants can also engage in non-verbal reasoning by exclusion, which may recruit negation (A or B, not A, therefore B). Although there is some controversy about whether this reasoning emerges around 17 months of age (Feiman, et al., 2022) or is already available to 12-month-olds (Cesana-Arlotti et al., 2018), there is consensus that it emerges significantly earlier than infants' mapping of negation to any negation word.

Still, there are alternative explanations of infants' successes on each of these non-verbal tasks, which leave the door open for the possibility that further conceptual change is still required to enable children to think negated thoughts. Infants' ability to learn a rule in which the last syllable is not the same as those before it (Hochmann & Toro, 2021) may be limited to previously seen syllables or sounds. 12- and 14-month-old infants' ability to reason by exclusion (Cesana-Arlotti et al., 2018, 2020) may be similarly limited to comparisons between what they see and information they have previously stored in visual (rather than auditory) working memory. Feiman et al. (2022)

argue that all existing evidence of pre-verbal negation (including additional evidence of their own) is currently compatible with infants having merely implicit representations of negation. Compared to full-fledged logical negation, these implicit representations can be limited in two independent ways. First, they might only negate certain kinds of content (e.g., information held in visual or auditory working memory; desires, as in Rejection uses; presence, as in Nonexistence uses), thus lacking the landmark content-independence and full compositional range of logical negation. Second, they might only be capable of some, but not all, of the logical functional roles of negation. In this sense, one kind of implicit negation is *contrariety*, which can determine that two inputs, like *red* and *blue* are mutually exclusive without one being the negation of the other (see Horn, 1989). For either type of implicit negation (domain-specific or functionally limited), it is possible that further conceptual development is necessary in order for the ability to think full-fledged negated thoughts to emerge, possibly using these pre-existing implicit representations as building blocks.

But just because alternative explanations of infants' successes can be given does not mean that those alternatives are right. It is also quite possible that young infants' successes do reflect their ability to think and reason with negation prior to their mastery of its linguistic expression. If not conceptual factors, what else could account for the documented delays in children's understanding of *no* and *not* as expressions of truth-functional negation? An alternative explanation is that the only major challenge children face is learning how their language expresses this concept.

1.3. Information Change: The case of missing evidence

It may be that what enables older children to learn that *no* and *not* express truth-functional negation is a changing ability to extract the relevant information from the language input children receive, independent of their ability to think negated thoughts. This explanation does not require that the input itself changes, only that the child's ability to extract information from it does. For example, if a child cannot yet segment the speechstream into words, it'd be harder to learn their meanings. Later on in acquisition, learning certain new words has been shown to require specific pre-existing linguistic knowledge. The parade case, studied extensively by Gleitman and colleagues, is learning verbs (Gleitman et al., 2005). Most verbs are extremely difficult to learn if one only has access to the co-present perceivable world, either because they are too abstract to ever perceive directly (e.g., *thinking*), because they are ambiguous in their construal (e.g., *chasing* events are also always *fleeing* events), or because the events being referred to are not present when the verb is uttered (e.g., almost all past tense uses of verbs). As Gleitman and colleagues showed, these verbs, especially ones with less imageable meanings, are delayed in children's vocabulary, compared to more imageable nouns. As proposed in Landau and Gleitman (1985), Gleitman (1990) and Gillette et al. (1999), the way children eventually learn the meanings of abstract verbs is by leveraging a growing understanding of how the surrounding linguistic context in which verbs appear constrains hypotheses about what they could mean. For example, a learner who sees an ambiguous event that can be construed either as a dog chasing a cat or a cat fleeing from a dog can tell which construal is meant by "The dog blicked the cat*",* if they know what *the dog* and *the cat* mean and also know that English syntactic subjects tend to be the agents of actions. Once a learner has access to the meaning of other words in the sentence and to the syntactic rules for their combination, they are equipped to tell which concepts verbs express, and to learn them rapidly (e.g., see Naigles, 1990; Yuan & Fisher, 2009; Nappa, Wessel, McEldoon, Gleitman & Trueswell, 2009). Critically, this view need not posit any conceptual change on the part of the learner to explain why some words are harder to learn than others. Infants could have a firm conceptual grasp of events, even abstract events, but identifying which word forms go with which concepts might require linguistic evidence that only a more advanced language learner has access to.

An information change account could likewise apply to the acquisition of negation words and to the progression observed by Bloom. Young infants may be able to think negated thoughts, and use negation in their own reasoning about the world. However, they might struggle to infer when others are expressing negation verbally if they cannot see how negation (or an entire negated phrase) could refer to anything observed in their immediate environment. In this case, inferring that negation is being used would require additional linguistic evidence, such as the syntax of the construction within which negation appears and the meaning of the words being negated. To see the idea, imagine hearing, "That cat is not nice" in the context of an angry cat. Compare how much easier it might be to infer that *not* expresses negation if you could already understand "That cat is nice," than if you were also simultaneously trying to infer whether *nice* means angry, or what a *cat* is.

Unlike truth-functional negation, other meanings of negators might be easier to infer without access to as much linguistic evidence. For example, when parents Reject something themselves or Prohibit their child from doing something, their use of a negator may be accompanied by various auditory cues (e.g., tone, prosody) and visual cues (e.g., finger wagging, being carried away, something being removed), cluing children into just those meanings. Nonexistence, in the form of *no* or *not,* may be accompanied by a characteristic hand-gesture (hands flipped with palms upwards to indicate uncertainty or ignorance (Vilà-Giménez, Dowling, Demir-Lira, Prieto & Goldin-Meadow 2021; Franklin, Giannakidou & Goldin-Meadow, 2011) or searching behavior. Thus, Nonexistence and Rejection might be produced earlier than Denials – the clearest expressions of truth-functional negation – simply because there are more substantial extra-linguistic cues to their meaning, and so, unlike Denial, they can be learned with less access to linguistic information.

Critically for our present work, the Information Change account makes a unique prediction: even conceptually equipped learners who are fully capable of thinking negated thoughts should nevertheless still struggle to infer that negation is expressed in Denial uses when they lack the relevant linguistic context, but should face less difficulty with Rejection or Nonexistence uses. If this prediction holds, then it would not be necessary to assume additional conceptual change to explain the young learner's Bloomian trajectory.

1.4. Human Simulations of Vocabulary Learning

To test this prediction, we use the Human Simulation Paradigm (HSP; Gillette et al., 1999). This method assesses the information that a conceptually equipped observer would need to infer that a parent is expressing different words in natural parent-child interactions. The HSP was first used to evaluate the information needed to infer the use of a particular verb (and, correspondingly, the expression of that verb's meaning), so as to examine how information availability may explain developmental changes in child vocabulary.

In the HSP, instead of instantiating a learning model in silico, human adults model the learner in vivo. They are given some partial input from naturalistic uses of a target word, and asked to guess what the parent said. Adults are inarguably conceptually equipped, so the HSP allows researchers to hold conceptual ability constant while manipulating information availability to "simulate" different stages of word learning. In one condition, designed to be analogous to the earliest stages of word learning when information is available only through visual observation, participants watch a short video of a natural parent-child interaction that is muted other than a single beep. They are tasked with guessing what the parent uttered at the beep (which was actually either a noun or a verb).

Gillette et al. (1999) found that when given only visual information, adults were much more likely to correctly guess the nouns that parents say than the verbs. On the other hand, in a condition where adults were given syntactic information in addition to the video, they were overwhelmingly more likely to guess the verbs. This accuracy is impressive given that in this condition, participants were provided with a 'jabberwocky' sentence (e.g., guess what GLORP means in "The dax GLORPS that zeb is flerpy"). These findings support an Information Change account and suggest that verbs are less frequent than nouns in early childhood production because verbs require understanding more linguistic context to learn – even for learners who are conceptually equipped. This original effect has been replicated (Snedeker, 1999; Piccin & Waxman, 2007; using child Point-of-View videos: Yurovsky, Smith & Yu., 2013) and the HSP has been extended to address other learning questions (Fitch, Arunachalam & Lieberman, 2021; Cartmill et al., 2010; Zhang, Yurovsky & Yu, 2015). Here, we extend it to investigate negation.

We present data from two independently conducted Human Simulation Paradigm experiments (one conducted at the University of Pennsylvania, the other at Brown University).[2] Both experiments compare different parental uses of negation words (for Prohibition, Nonexistence, and Denial), and test whether Denial uses – in which the negator can express only truth-functional negation – are especially difficult to identify without additional linguistic information. The experiments differ in many other respects: whether the adult participants saw videos from a third-person view or a child's first-person perspective, whether they had to guess an entire utterance or only one omitted word, whether linguistic information (when it was given) appeared after the video ended or concurrently, among others. On the one hand, any of these methodological differences could lead to diverging results. On the other hand, these differences suggest that any converging results are robust, as they would be unlikely to stem from particular features of either design. The logic is the same in both experiments. Restricting the language input of conceptually mature adults to only visual information simulates the earliest stages of word learning. Loosening the restriction to allow more linguistic input simulates the information available at later stages of word learning. If, under these simulated conditions, adult participants recapitulate the learning trajectory of children – identifying Prohibition and Nonexistence uses from visual input alone, but identifying Denial uses only when more language input is available – then, without excluding the possibility of additional Conceptual Change) Information Change alone would be sufficient to explain why only older children produce Denials and understand *no* and *not* as expressing truth-functional negation.

1. Experiment 1

The experiment method and sample size was preregistered at: https://osf.io/76rhk/.

2.1. Methods

2.1.1. Participants

Participants were 257 adults recruited either from Prolific (N=165) or from the University of Pennsylvania's undergraduate participant pool (N=92) and participated online. Each participant was assigned to one of three between-subject conditions (of Information Type, see below) resulting in 72 participants in each. This number was chosen based on a power analysis of Experiment 3 in Gomes, Huh, and Trueswell (2020). Participants from the University of Pennsylvania had the option of receiving

[2] Experiment 1 was conducted by authors V.G. and J.T. at the University of Pennsylvania, while Experiment 2 was conducted separately by authors R.F., R.D., and D.S. at Brown University. These groups knew about each other's study but did not attempt to match the studies' design or execution.

class credit. Otherwise, those in the Video-Only and Video-and-Language conditions were paid $6.50. The Language-Only condition was shorter in duration, and participants were paid $4.00.

An additional 31 adults participated but were excluded from analysis because at least 40% of their responses showed clear signs of inattention (primarily due to providing responses containing more than one word in the Language conditions). This eliminated 2 participants in the Language-Only, 3 in the Video-Only, and 27 in the Video-and-Language condition. Our preregistered exclusion criteria also included unusually fast reaction times, repetitive responses (e.g., always saying the same thing) responding with nonsensical or ungrammatical sentences, but no participants were excluded for these reasons. Finally, 10 other participants were excluded for self-reporting being non-native English speakers.

2.1.2. Materials

2.1.2.1. Database

The video stimuli for this experiment came from the Language Development Project Corpus, which includes video from a third-person view, audio, and tagged transcripts of parent-child interactions in 64 families with typically-developing children from the Chicago area that were recruited to provide a representative sample of the area based on the 2000 census (Goldin-Meadow et al., 2014). The transcripts in the corpus are tagged to include original utterances, morphology, parts of speech, and gesture of both the primary caregiver and the child. To build this corpus, researchers visited the parent and child in the home every four months for 90 minutes to record the parent and the child interacting in everyday activities (e.g., playing, eating, cleaning up, reading, and completing other work around the house). This began when the child was 14 months of age and continued until they were 58 months of age.

2.1.2.2. Sampling Procedure for Stimuli

We randomly selected four parent-child dyads for stimuli creation for our study. The only selection criteria were that two children had to be male and two female, and that two had to be high-SES and two low-SES (such that there’d be one of each; determined as in Cartmill et al., 2013). From these four dyads, we extracted all parental utterances (n=9971) from the first three home visits (14, 18, and 22 months of child age). These ages were selected because they are all prior to when children typically begin to regularly produce Denials (Bloom, 1970; Pea, 1980; Choi, 1988; Nordmeyer & Frank, 2018a), and before they begin comprehending Denial on tasks that test exclusion inferences (Feiman et al., 2017; Austin et al., 2014). This procedure minimized any possible variation in the input that could stem from parents tailoring their productions to children's emerging understanding of negation. A total of 790 utterances were determined to contain negation (i.e., they included *no*, *not*, *-n't*) and were subjected to further annotation. The mean length of utterances (MLU) for those utterances which included a negator was 3.2, and 2.7 for those which did not.

2.1.2.3. Video Coding

The first author (VG) and two Research Assistants annotated the parents' negative utterances to code for Negation Function. Each negative utterance was coded according to whether it corresponded to Denial, Nonexistence, and Prohibition. Bloom (1970), who examined children's utterances rather than parents' utterances, annotated negation uses of Denial, Nonexistence, Rejection (instead of Prohibition). Choi (1988), who also coded only children's utterances, expanded Bloom's taxonomy with additional functions of negation, including Prohibition. As Choi (1988) notes, Rejection and Prohibition largely differ based on the social hierarchy that an interaction takes place within (parents typically prohibit their children, not vice versa; children are usually in a position to reject what parents are trying to make them do, not vice versa). Parents therefore produce Prohibitions much more than Rejections. Since parental rejections are rare, we used a merged Bloom-Choi taxonomy of negation for our analyses of parental utterances: Denial, Nonexistence, and Prohibition (Bloom, 1970; Choi, 1988). The first author instructed the two other coders on how to annotate videos (see supplemental materials). Each coder was then given 20 practice videos to annotate and received feedback from VG after they were done. Once it was determined the coders understood the annotation procedure, they each annotated one-third of the dataset separately. After this process was complete, the coders completed a final set of 20 items which were held out to assess inter-rater reliability. There was unanimous agreement on 85% of cases and majority agreement on the remaining 15%.

2.1.2.4. Stimulus Construction

2.1.2.4.1. Visual materials

Eighteen videos of parent-child interactions involving negation were randomly selected from our annotated subsample (six from each type of negation: Prohibition, Nonexistence, Denial). To be selected, videos had to meet the following criteria: 1) both the parent and the child must have been visible during the critical utterance, 2) no private or sensitive information was visible in the videos, 3) the target utterances were not directed at siblings or others, and 4) the target utterances did not come from the parent reading.

For each negative utterance, a corresponding affirmative utterance was randomly selected from the corpus based on the prejacent that was negated (e.g., if a negative utterance was "That's not your ball," a matched affirmative utterance would include *ball*).

Eighteen additional filler utterances were randomly selected from utterances that did not contain a negator. All of these videos (18 affirmative, 18 fillers) had to meet the same inclusion criteria, above. These stimuli were randomly divided into two lists of 27 videos each, and each participant saw one list. For analysis, participants were randomly paired with other participants of the complementary set to form a pair. As specified in the pre-registration, analyses were conducted on the pairs.

For each utterance, the video of the interaction was extracted and edited to be approximately 40s using DaVinci Resolve, with the target utterance appearing at 30s, and the video ending 10s after the onset of the utterance. Each video was then muted, and a tone was added where the critical utterance occurred in the video (i.e., contiguous with the entire critical utterance).

2.1.2.4.2. Linguistic materials

In the conditions with linguistic information (Language-Only, Video-and-Language), a modified transcription of the target utterance appeared after the video. For negated utterances, the negator(s) was removed and replaced with a blank. Contracted items were removed in full (e.g., “That isn’t him,” would become “That ___ him.”). For affirmative utterances, a blank was inserted where negation would have occurred had it been similarly negated (e.g., “The doggie is yellow!” would become “The doggie is ___ yellow!”). For each filler item, a verb or noun was randomly chosen to be replaced by a blank. For example, “The dog ran home,” became either “The ___ ran home,” “The dog ___ home,” or “The dog ran ___”).

2.1.3. Procedure

Participants completed the task online through the PCIbex platform (Zehr & Schwarz, 2018). Participants in the Video-Only and Video-and-Language conditions viewed one of two lists of 27 videos of parent-child interactions. To ensure participants in these conditions would be able to hear the tone, they were asked to enter a numeric code played via audio, and upon entering it correctly they could begin the experiment. Participants were told that they would be watching videos of parents interacting with their children, and that the video would be muted except for one continuous tone. They were told that their task was to guess what the parent had said during this tone. In nine of the videos, a negator had been uttered by the parent. Out of the remaining 18 videos, nine were matched affirmatives, and nine were fillers. In the Video-Only condition, participants saw only the videos and were not presented with additional information after it played. They were asked to guess the entire utterance spoken at the time of the tone. In the Video-and-Language condition, after each video finished playing, participants were presented with additional linguistic stimuli from the parents’ actual utterance (described in 2.1.2.4.2. above). These stimuli had a missing word (indicated by a blank), and, instead of being asked to guess the whole utterance, participants were asked to only guess what the missing word was. As indicated above, the blank for the matched affirmative utterances appeared where the negator would have gone if the sentence had actually been negated. Thus the blank in this condition did not indicate a missing word, although the participants were led to believe it did.

Information Type	Final screen for a Negative trial
Video-Only Condition	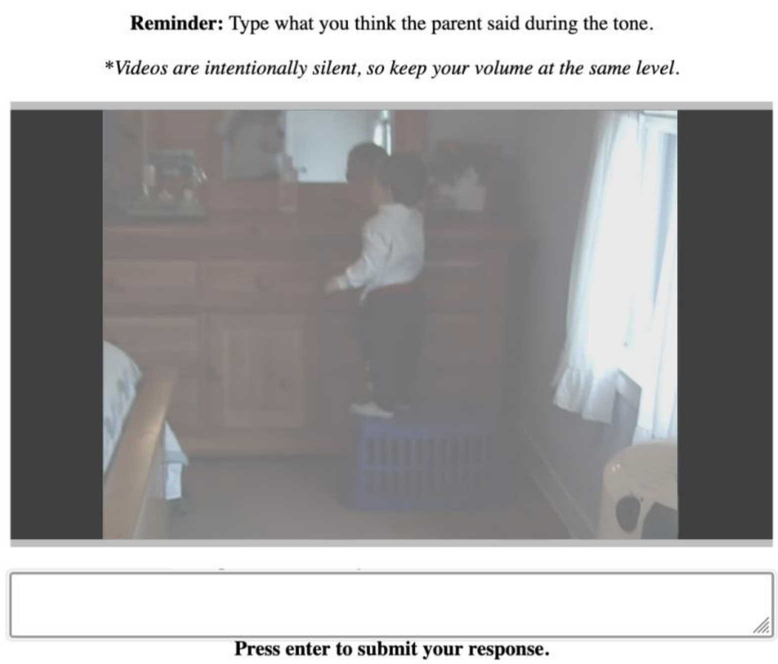
Video-and-Language Condition	

	Reminder: Type what you think the parent said during the tone. **Videos are intentionally silent, so keep your volume at the same level.* 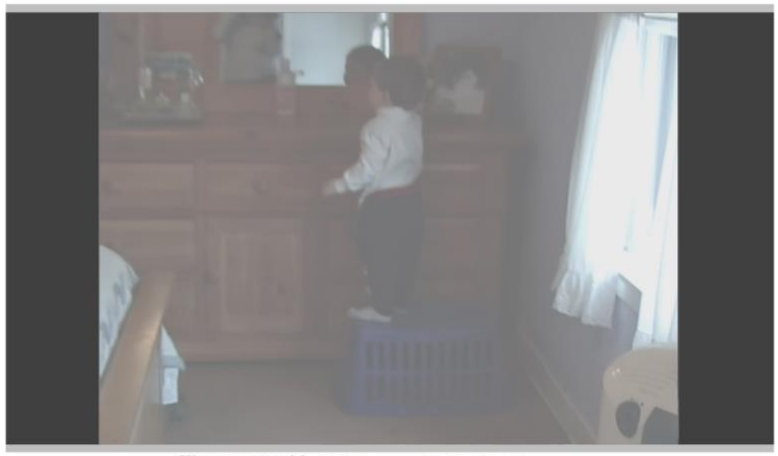 The parent said: "Johnny, ___ that's glass." **Press enter to submit your response.**
Language condition	The parent said: "Johnny, ___ that's glass." **Press enter to submit your response.**

Table 1: The final screen of a trial shown in each Information condition. For the conditions that included Video, the textbox (and the linguistic stimuli in the Video & Language condition) only appeared after the video finished playing. After the video played once, it would not play again.

Participants in the Language Only condition did not view any videos, but saw the same linguistic stimuli presented to participants in the Video-and-Language condition and were similarly asked to guess the missing word.

2.2. Analysis

2.2.1. Variable coding

In both sets of analyses below, the dependent variable was the presence of a negator (*no*, *not*, or *-n't*) in participant responses, coded as a 1 (includes negator) or 0 (does not include negator). Coding schemes for the independent variables were designed to test the most theoretically interesting comparisons. Analyses that involve comparing pairs of Information Types are performed on the subset of the data in those conditions, which are then sum-coded (as either Language-Only vs. Video-and-Language, -1 and +1, or Video-Only vs. Video-and-Language, -1 and +1). For comparison of negative utterances and matched affirmatives, Utterance Polarity is similarly sum-coded (Affirmative vs. Negative, sum-coded as -1 and +1). For comparison of types of negation, Negation Function is Helmert coded to allow for two comparisons of interest: Nonexistence-vs-Prohibition, which excluded data from the Denial condition; and Denial-vs-Other, which compared Denial to Nonexistence and Prohibition collapsed together.

2.2.2. Statistical modeling

For all analyses, we tested our predictions using multilevel logistic regression models on trial-level data (glmer from lme4 in R; Bates et al., 2015) which permit statistical inferences across participants and items simultaneously. In these and all other multilevel models reported in this paper, we tested the significance of fixed effects by using likelihood ratio tests on the Chi-square values from nested model comparisons with the same random effects structures (using the car package in R; Fox & Weisberg, 2011). For each model, we began with the maximum random effects structure given the experimental design (e.g., an intercept and effect of Utterance Polarity for participants, and an intercept, effects of Information Type, Utterance Polarity, and their interaction for Item). In cases where models did not converge, we simplified the random effect structure by removing the random slope which accounted for the least variance.

For all analyses, we report effect sizes in terms of odds ratios (ORs) by taking the absolute value of the exponent of β (the logit-transformed fixed effect coefficient). Thus, the OR represents the increased likelihood of a participant producing negation in their response for one condition vs. another. For example, an OR of 4.0 for Utterance Polarity (Affirmative vs. Negative) would mean that a participant's response was 4 times more likely to contain negation when the parent had produced negation, compared to when the parent had not produced negation. We then derived p-values for each reported effect from likelihood ratio tests comparing models with and without that effect.

As the Information Type comparisons of interest are between each of the partial information conditions (Video-Only, Language-Only) and the full information condition (Video-and-Language), we compared these pairs of conditions in separate analyses. Similarly, all simple effects were conducted by subsetting the data and running a separate model on the conditions of interest.

For the first set of analyses, we consider the difference between negative utterance vignettes and their matched affirmative vignettes (i.e., excluding fillers), and how this difference is affected by the type of information available to participants. These models predict the presence of a negator in participants' responses using Utterance Polarity (Affirmative vs. Negative), Information Type (Video-Only vs. Video-and-Language; Language-Only vs. Video-and-Language) and their interaction as fixed effects. In the second set of analyses, we exclusively consider performance in negative utterance vignettes (i.e., excluding matched affirmatives and fillers). Here, we are interested in whether participants' ability to tell what parents said depends on the

Negation Function, and its interaction with the type of information available to them. These models predict the presence of a negator (*no*, *not*, or *-n't*) in participant responses using Information Type and Negation Function (Denial, Prohibition, Nonexistence) and their interaction as fixed effects.

2.3. Results

2.3.1. Negation overall

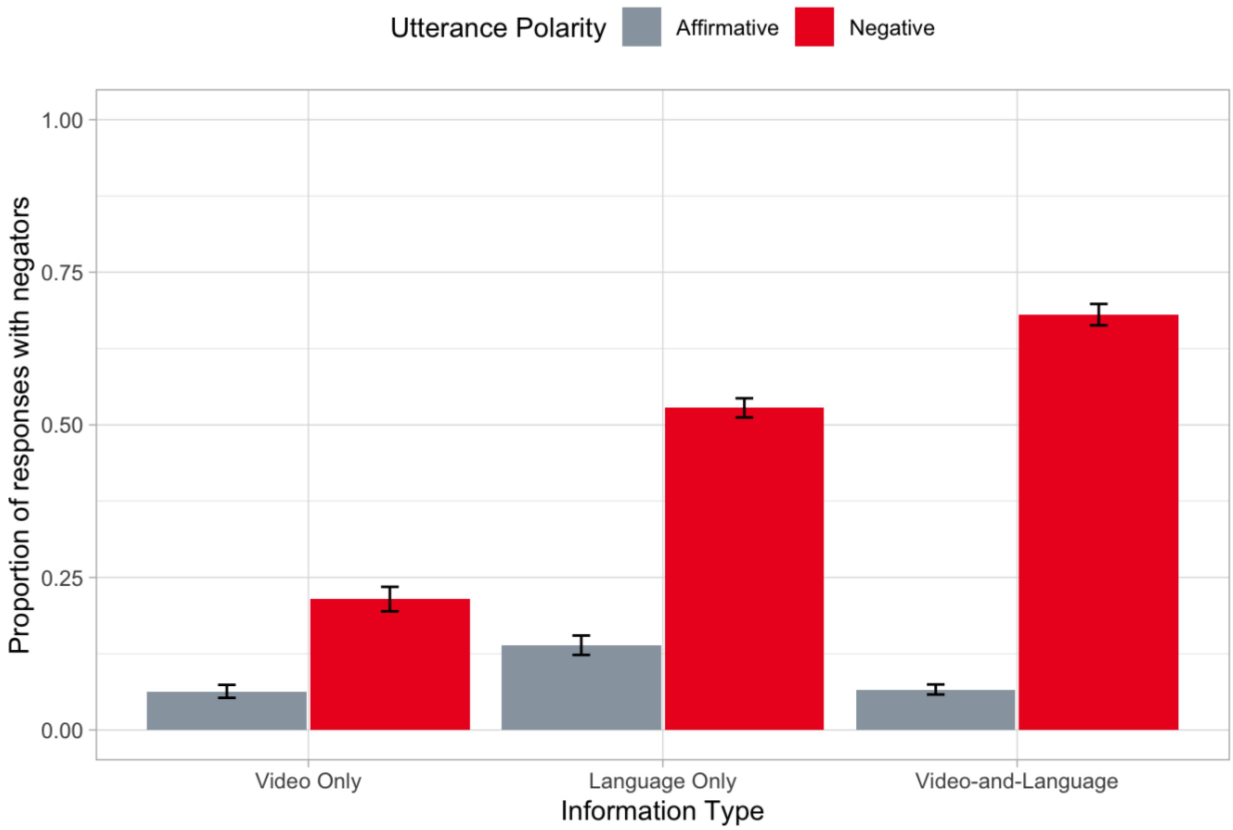

Figure 1: Proportion of negators included in participant responses for all Information Type conditions. Data are presented by Utterance Polarity (Negative, Matched Affirmative) and the amount of information provided along the x-axis (Information Type) along the x-axis. Error bars indicate ±1 Standard Error.

Figure 1 shows the proportion of participant responses that contained a negator (i.e., *no, not, -n't*) as a function of the information provided to participants and the kind of parental utterance they had seen. To investigate the benefit of receiving linguistic information in addition to the visual scene, performance in the Video-Only condition was compared to performance in the Video-and-Language condition. There was a significant main effect of Utterance Polarity (β=1.43, z=6.53, $p<0.001$, OR=4.18), showing that participants were much more likely to include a negator in their response when the parent had actually used a negator in that utterance (Negation condition: 45%), compared to when the parent had not (Matched Affirmative condition: 6%). There was additionally a main effect of Information Type, such that participants were more likely to provide a negator overall in the Language-and-Video condition (37%) compared to the Video-Only condition (13%) (β=0.68, z= 7.99 p<0.001, OR=1.979). Additionally, there was a significant interaction between these two factors (β=0.65, z=9.24 p<0.001, OR=1.91), which reflected the fact that the effect of Utterance Polarity (Negative vs. Affirmative) was larger in the Video-and-Language condition than in the Video-Only condition. Simple effects analyses revealed that the effect of Utterance Polarity (Negative vs. Affirmative) was significant in both the Video Only, (β=0.925, z=3.41 p<0.001, OR=2.52), and the Video-and-Language, $\chi^2 = 38.17$, $p < 0.001$ (β=2.70, z=9.24 p<0.001, OR=14.83) conditions.

To investigate the added benefit of seeing the visual scene compared to linguistic information alone, performance with Language-Only was also compared to performance in Video-and-Language. An analysis of these data similarly showed a significant main effect of Utterance Polarity (β=1.90, z=5.70 p<0.001, OR=6.67). Across these two conditions, participants were again more likely to provide a negator when the parent had actually used a negator (Negation condition: 60%), compared to when the parent had not (Affirmative condition: 10%). There was also a marginal effect of Information Type, (β=0.09, z=1.115 p = 0.26, OR=1.098), with participants marginally more likely to provide a negator in the Language-and-Video condition (37%) than in the Language-Only condition (33%). Importantly, there was a significant interaction between these two factors, reflecting the larger effect of Utterance Polarity (Negative vs. Affirmative) in the Video-and-Language condition than in the Language-Only condition (β=0.55, z=8.36, p<0.001, OR=1.73). As in the Video-and-Language condition (reported above), the simple effect of Utterance Polarity (Negative vs. Affirmative) was also significant in the Language Only condition (β=1.41, z=4.14, p<0.001, OR=4.08).

2.3.2. Negation by type

As noted in the Methods, participants saw an equal number of each of the three different types of negation: Denial (e.g., “That’s not a rock”), Nonexistence (e.g., “No doggies”), and Prohibition (e.g., “Don’t do that”). Here we examine whether some of these types were easier to guess than others, and whether they benefitted differentially from different kinds of Information. Figure 2 shows the proportion of participant responses that contained a negator (any of *no*, *not*, or *-n’t*) when presented with negative parental utterances that really did include a negator, as a function of the information provided to participants and the type of Negation used (Denial, Nonexistence, or Prohibition).

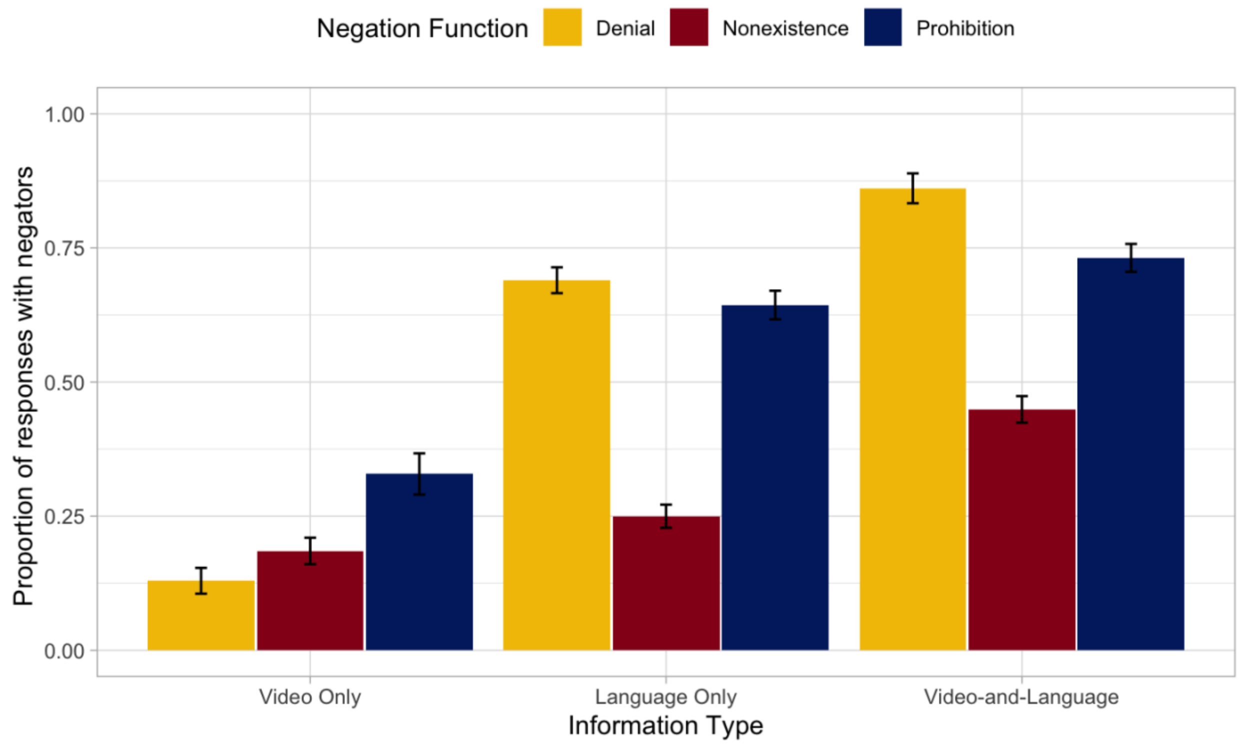

Figure 2: Proportion of negators included in participant responses for vignettes in which the target utterances included a negator by communicative function of that word (Denial, Nonexistence, Prohibition) and the amount of information provided (Information Type) along the x-axis. Error bars indicate ±1 Standard Error.

As Fig. 2 shows, linguistic information greatly increased participants' likelihood of responding with a negator when the vignette showed a Denial use of negation. To analyze this effect, a regression modeled the effect of Information Type (as performance in the Video-Only condition compared to performance in the Video-and-Language condition), the effect of Negation Function, and their interaction. The main effect of Negation Function was a marginally significant predictor (χ^2 (2)= 5.92, p = 0.052). Examining the Helmert-coded contrasts of each Negation Function within this model, Prohibition vignettes elicited significantly more negator responses than Nonexistence vignettes (β=1.297, z=2.11, p<0.05, OR=3.66), whereas Denial vignettes did not result in significantly more negator responses than the average of Prohibition and Nonexistence vignettes (β=0.47, z=0.86, p=0.391, OR=1.6). Importantly, however, there was a significant interaction between the effects of Negation Function and Information Type (χ^2 (2)= 41.55, p < 0.001), with Denial vignettes benefiting more from linguistic information than Prohibition and Nonexistence vignettes (β=1.17, z=6.33, p<0.001, OR=3.22). Prohibition vignettes, on the other hand, did not benefit from linguistic evidence any more than Nonexistence vignettes did (β=0.24, z=1.36, p=0.17, OR=1.27). As before, participants were more likely to produce negators when given both Video-and-Language, compared to Video-Only (β=1.34, z=13.42, p<0.001, OR=3.82). Simple effects analyses revealed that the effect of Information Type was significant for all negation functions (Denial: β=2.75, z=4.398, $p<0.001$, OR=15.69; Prohibition: β=1.82, z=2.37, $p<0.05$, OR=6.17; Nonexistence: β=1.06, z=2.37, $p<0.05$, OR=2.88)

To investigate the benefit of visual information, performance in the Language-Only condition was compared to the Video-and-Language condition. The effect of Information Type had a significant impact on participants' performance (β=0.66, z=6.28, $p<0.005$, OR=1.93). Once again, Negation Function was Helmert-coded to compare Denial to Nonexistence and Prohibition (collapsing the two) as well as Nonexistence to Prohibition. Within this model, Denial utterances were marginally more likely to elicit a negator response (78%) than Prohibition (35%) and Nonexistence utterances (69%) in this comparison, χ^2 (2)= 6.50, $p < 0.05$ (β=1.92, z=1.72, p=0.085, OR=6.795), and Prohibition similarly resulted in more negator responses than Nonexistence (β=2.46, z=1.27, p=0.052, OR=11.68). Importantly, however, there was no significant interaction (χ^2 (2) = 4.97, p = 0.08). Like the Video-and-Language condition, and unlike the Video-only condition, participants in the Language-only condition were more likely to respond with a negator to Denials than to the other two Negation Functions. In short, our prediction was borne out, that Denial would benefit from the addition of language more than Prohibition and Nonexistence. Moreover, this effect is not simply due to Denial benefiting more from *any* additional information, since only the addition of language information – not the addition of video to language – interacted with Negation Function to affect participants' responses.

We conducted a post-hoc analysis to compare negative items with their matched affirmatives (Supp. B). For these analyses, affirmatives were designated as the same function as the negative utterances they were derived from (e.g., Denials if they had originally been derived from a Denial utterance). Critically this analysis focuses on the Language Only and Video-and-Language conditions, since the linguistic information paired with the matched affirmatives was artificially manipulated so as to allow for the addition of a negator. Participants were more likely to include negators in their responses for real negative items than in the corresponding affirmative controls (Prohibition (65% vs 15%;β=1.98, z=0.52, $p<0.001$, OR=1.24), Nonexistence (39% vs 7%;β=2.55, z=0.51, $p<0.001$, OR=12.85) and Denial (78% vs 9%; β=1.92, z=1.72, p=0.085, OR=6.795)). While Nonexistence and Denial pairs benefitted overall from having both Video and Language as indicated by the significant effect of Information Type (Nonexistence: β=0.2998, z=0.12, $p<0.05$, OR=1.28; Denial: β=0.25, z=.12, $p<0.05$, OR=6.795), Prohibition did not (β=-0.22, z=0.14, p=0.11, OR=7.25). However, all had significant interactions between Polarity and Information (Prohibition: χ^2 (1)= 30.396, $p<0.001$ (β=0.62, z=0.11, $p<0.001$, OR=1.86); Nonexistence: χ^2 (1)= 14.43, $p < 0.001$ (β=0.45, z=0.12, $p<0.001$, OR=1.57); Denial: χ^2 (1)= 18.58, $p < 0.001$ (β=0.5, z=4.31, $p<0.001$, OR=1.65)). For Prohibitions, however, this seems to be because the addition of Video decreases negator guesses in the matched affirmatives but does not significantly affect the negative utterances. For Denials and Nonexistence the interaction suggests that where the negative utterances benefitted from the addition of visual information, the matched affirmatives did not, even though linguistically they permitted the addition of a negator.

2.4. Discussion

We found support for the key predictions of the Information Change account, namely, that 1) with only video and without additional linguistic information, it is especially hard to infer a negator when parents are using it to deny, and 2) additional linguistic information helps to infer a negator when it was used to deny more than for any other uses of negation.

Thus, subjected to different restrictions on the availability of language information (that is, simulating learners who can access different kinds of information), conceptually equipped adults recapitulated the developmental trajectory seen in child production. When limited solely to observation of the world like the youngest language learners, the adults were unlikely to infer that a parent was expressing truth-functional negation. When they did correctly infer that a negator was used, it was usually when it was used to Prohibit or comment on Nonexistence, but rarely when it was used in its truth-functional sense, to Deny. However, with linguistic information made available, adults were much more likely to correctly infer the presence of a negator when it was used to Deny, doing so 78% of the time. This suggests that the relatively late mapping of negation to negators seen in children's acquisition (both in comprehension and in production) can be explained by the changing information available to the learner who has to map truth-functional negation to a negator. Though this does not rule out the possibility of additional conceptual change, there is no need to invoke development either of the concept of negation itself, or of any other supporting cognitive abilities to explain the same trajectory in development. If children are unlikely to infer that negation is being expressed until they understand more of the rest of the language they hear, they will learn how to express negation relatively late – only after they know more of that surrounding language.

Interestingly, while the addition of linguistic information encouraged negator guesses when the target utterance did include negation, it did not increase negator guessing in the matched affirmatives, despite being grammatically permitted within the sentence. In the Language Only condition, participants were less likely to guess a negator for matched affirmatives, suggesting the role of pragmatic or semantic factors in guessing negators (e.g., participants assuming parents would be less likely to stop children from eating vegetables than rocks). However, the fact that the addition of video information increased negator guesses for Negative items and not Affirmative suggests an additional contribution of the scene. We will return to questions about the possible information present in scenes in the General Discussion.

In Experiment 2, we report on another HSP study that was conducted independently but in parallel, and which replicates and extends these findings. The two experiments made different but complementary design decisions, which is important for two reasons: first, because any converging patterns of results between them are less likely to be due to features of the design that are specific to one or the other experiment but are not theoretically central; second, because the specific features of each design enable different inferences.

For example, while Experiment 1 presented participants with the same number of vignettes of each Negation Function, the negator words used in these vignettes were unbalanced, reflecting their natural distribution of their communicative uses. As a consequence, no analyses by negator (e.g. *no* vs *not*) were presented here, since Negator was naturally confounded with Negation Function. As it turns out, analyses of *no* would never include any Denials, and analysis of *not* would never include Prohibition or Nonexistence. Indeed, in the Denial vignettes, *not* and *-n't* each appeared 3 times with *no* never appearing, while in the Prohibition and Nonexistence vignettes combined, *no* was used 7 times, *-n't* was used 5 times, and *not* never occurred. A consequence of this distribution is that it makes it possible that the differences in participants' ability to guess that a negator had been used when seeing vignettes of different Negation Functions were driven not by those functions, but by the differential identifiability of different negators in different contexts. Maybe *no* is easier to guess than *not* given only visual information, whereas *not* is easier to guess given the rest of the utterance. Indeed, because *not* appears only as an adverb modifying auxiliary verbs, whereas *no* can be a quantifier, a single-word utterance, or an anaphor (inter alia), it is plausible that the benefit of language information for identifying Denials could reflect the relatively higher predictability of *not* from its surrounding syntactic frame.

Experiment 2 tests this alternative by separating out Negation Function from Negation Word and crossing these factors in its design, thereby balancing the appearance of combinations of these factors with each other and with Information Type. However, this design has a corresponding downside – it is possible that the naturally occurring correlation between form and function is exactly what makes certain utterances easier to identify. If, for instance, Denial-*no* is harder to infer than Denial-*not* just because it is a rarely occurring, less typical usage, then manipulating form and function as orthogonal factors could underestimate participants' ability to infer the presence of a negation in more typical instances. The clearest picture should emerge from triangulating between the data patterns obtained from both experiments to identify what is consistent between them.

2. Experiment 2

Although Experiment 2 was similar in logic and approach to Experiment 1, there were a number of methodological differences between them, with different advantages and disadvantages with respect to the inferences they license. We summarize these here.

First, as highlighted above, the experiments differed in whether they manipulated the negator separately from the communicative function it expressed. The vignettes of Experiment 1 were randomly sampled from the corpus by Negation Function, reflecting the natural distribution of what each negation word is used to express. In contrast, Experiment 2 separated Negation Word from Negation Function, balancing the number of vignettes in which Denial and Nonexistence were expressed with *no* and with *not*. This allowed testing whether Denial uses are easier to guess with additional language information, independently of whether they were expressed with *no* or *not*. Relatedly, participants in Experiment 2 were always asked to guess the specific target word (rather than the entire utterance, as in the Video-Only condition of Experiment 1), and were credited with correct guesses only if they guessed the exact word that the parent had said (e.g., guesses of *no* when the target word was *not* would be coded as incorrect; though we also report exploratory analyses crediting participants for using any negator, below). Unlike Experiment 1, Experiment 2 did not include any instances of *-n't* as target words.

The experiments also differed in the kinds of videos they used. While Experiment 1 used videos from a 3rd person perspective, Experiment 2 used videos taken from a child's point of view. This choice was motivated by the possibility that a child's view might be a more informative source of visual evidence (see, e.g., Yurovsky, Smith & Yu, 2013, though also see Gleitman & Trueswell, 2020). If the visual information in the child's view is good enough to identify Denials without any language input, we would expect no difference between the Video-Only and Video-and-Language conditions (contra Experiment 1). On the other hand, if the difference found in Experiment 1 replicates in Experiment 2, it would suggest that Denial negations are harder to identify without language input, across a wider variety of visual information (both 1st- and 3rd-person). Relatedly, the differences between corpora resulted in all items being constructed from sampling from a single parent-child dyad, unlike Experiment 1 which sampled from four dyads.

When presenting language information, the two experiments differed in how much language they included. In Experiment 1, like most prior HSP studies, conditions that showed language information presented participants with just the one utterance in which the target word appeared and no language from the rest of the video. This utterance appeared after the video ended, around 10 seconds after the target word had been said, like the intertitle text shown in silent films. This conservative manipulation tests for effects of relatively minimal language input, simulating understanding only the target utterance, separated in time from its visual context. In contrast to this approach, Experiment 2 aimed to make it easier for adult participants to simulate learners who might have full access to the type of information presented in their assigned condition. Thus, in Experiment 2, language information was provided through contemporaneous utterance-length subtitles that spanned all of the speech in the entire 40s video, appearing both before and after the target word (with any additional instances of the target word also beeped out and excluded from the subtitles). Additionally, because the video was subtitled as it unfolded, more information about the timing of all speech relative to the visual environment was available to viewers. This aimed to allow participants to simulate a learner that could benefit from the language information in the discourse context surrounding the target word and utterance.

The two experiments also differed in their control targets – the kinds of utterances and words they asked participants to guess for comparison to the negators. In Experiment 1, participants were presented with matched affirmative utterances, selected from the corpus to mention the same prejacent as the negatives. This allows comparing the guessability of utterances with similar content, but requires participants to guess a word that was not actually said in that utterance (marked by a blank inserted where a negation could grammatically fit in the Language-Only and Video-and-Language conditions). Experiment 2 instead compared the identifiability of *no* and *not* to the identifiability of the 20 most frequent nouns, verbs, and function words. Although this makes vignettes with target negators less comparable to other trials in terms of their content, it situates negation within the noun-verb comparisons made in earlier HSPs (e.g., Gillette, et al.,1999; Piccin & Waxman, 2007; Yurovsky, et al., 2013) and extends those comparisons to other function words. The noun, verb, and function word target trials also helpfully act as fillers, so that the correct target word would not always be a negator.

Finally, Experiment 2 explored the effects of additional Information Types, following Gillette, et al. (1999). In addition to conditions in which participants saw either none or all of the words surrounding the target (corresponding to the Video-Only and Video-and-Language conditions in Experiment 1), Experiment 2 included three other conditions: one in which only the content words were subtitled throughout the video and all function words were masked, one in which only the function words were subtitled and content words were masked, and one in which content words were subtitled but arranged in an alphabetized list and separated by commas, removing word order information (see Table 2). These conditions aimed to test whether, in guessing that a negation word had been used, participants benefited more from information about the prejacent content being negated or from information about the syntax of the utterance. In contrast to Experiment 1, Experiment 2 did not include a Language-only condition, limiting the ability to make inferences about the independence of visual and language information.

3.1. Methods

3.1.1. Participants

540 participants who self-reported being monolingual English speakers and at least 18 years old were recruited through Amazon Mechanical Turk with the aid of TurkPrime (Litman, Robinson, & Abberbock, 2017). Participants were compensated $3.50 upon completion of the study. An additional three participants were excluded because they did not complete the entire study.

3.1.2. Materials

The stimuli were created in several steps. Taking a video corpus of parent-child interactions, we first used automatic machine transcription to identify the most frequent nouns, verbs, and function words spoken by parents, as well as all parental uses of *no* and *not*. We then manually coded these according to Negation Function, as in Experiment 1. From the resulting database, we sampled video clips in which parents said specific words in order to create 4 distinct lists, containing unique nouns, verbs, and function words, as well as different clips of each Negation Function-and-Word combination. Finally, five versions of each clip were created to simulate five different conditions of information availability (Information Type). Each step is detailed below.

3.1.2.1 Database

Video stimuli were drawn from Sullivan et al.'s (2020) SAYCam corpus, which is comprised of recordings of parent-child interactions seen from the child's view, filmed with a head-mounted camera. We used all available videos from one child, Asa, because his recordings had the best audio quality and were the only ones on which our machine transcription methods worked reliably (see below). Recordings of Asa were made at least twice per week for 1-1.5 hours in each session, between 7 and 23 months of age. This produced 314 videos, totalling over 138 hours.

3.1.2.2. Sampling Procedure for Stimuli

Google's speech2text API was used for initial transcription of all footage. From this transcription, we identified the 20 most frequent words produced by Asa's parents in each category -- Nouns, Verbs, and Function Words -- to serve as targets.[3] A human coder then found all instances of these 60 words, as well as of *no* and *not*, and checked the machine transcription for the target words and surrounding speech, correcting as needed. The coder also made sure that the target word appeared as the intended part of speech. For instance, “have” and “do” frequently appeared as both main and auxiliary verbs, and coders separated clips in which these were the target words into the Verb and Function Word categories of stimuli, respectively.

To be included in the stimulus set, videos then had to meet the following criteria: the target word must have been part of an utterance directed at the child (not adults speaking to each other), must not be identifiable by lip-reading a visible speaker or reading text (from a book or any object visible in the video), and could not be part of a song (either one played in the background or sung by an adult). We also excluded videos that were very dark or blurry, but made no other attempt to ensure that any particular objects or referents of speech were present in the video (cf. Yurovsky, Smith, & Yu, 2013). The videos that fit these selection criteria contained, on average, 802 instances of each target word from the Nouns, Verbs, and Function Words (ranging from 78 instances of “chicken” to 5473 of “you”). From this set, we randomly chose 5 instances of each of the 60 target words.

We also identified 249 instances of *no* and 58 instances of *not*. These were coded by two coders, working independently and using the same coding scheme as in Experiment 1 (see Supplemental Materials) to categorize negation as Denial, Nonexistence, or Prohibition. Additionally, we added Rejection as a separate category, drawn from Bloom’s (1970) coding scheme. Prior to coding these instances, both coders were first trained on transcripts outside of the stimulus set, drawn from the CHILDES database, until they reached over 90% agreement with each other. In their subsequent coding of the stimulus set, they disagreed with each other on 44 instances of *no* and 13 instances of *not*, which were excluded. Of the remaining instances, the coders agreed on 176 instances of Prohibition *no*, 3 instances of Prohibition *not*, 18 instances of Denial *no*, 42 instances of Denial *not*, and 10 instances of Nonexistence *no*. The coders only agreed on one instance of Rejection *no*, and no instances of Rejection *not* or Nonexistence *not*, all of which we had planned to include as stimuli had we been able to identify enough of them. Although Rejection is the most common function of negation in children’s early speech (Bloom, 1970; Pea, 1980; Choi, 1988; Nordmeyer & Frank, 2018a), our findings are consistent with prior corpus evidence (Cameron-Faulkner et al., 2007) that Rejections are relatively rare in parents’ speech, and that Prohibitions are very frequent. Of the instances on which coders agreed, we randomly chose 5 each for Denial *no* and *not*, and for Nonexistence and Prohibition *no*.

[3] The initial transcription included “job” and “weather” among the 20 most frequent nouns spoken by parents. However, human coders determined that a high proportion of these instances were errors of the machine transcription (probably due to high priors on these words in Google’s API). We decided to discard these two words and replace them with the next two most frequent to maintain a set of 20 nouns.

For Prohibition *not*, we used all 3 of the instances on which both coders agreed. Thus, we included 5 combinations of Negation word by Type.

Finally, the target Negations, Nouns, Verbs and Function Words were divided into four Lists, with each List containing five of each type of word, for 20 words in total. For each word, each video clip instance in which that word was the target was assigned a Clip Number from one to five (except for Prohibition *not*). This produced five different sets of videos for each list of words.

3.1.2.3. Stimulus Creation

3.1.2.3.1. Visual Materials

Each video was trimmed down to a clip that started 30 seconds before the onset and ended 10 seconds after the offset of the target Word. Each video was then muted, and a tone was added at the time when the target word occurred in the video, using Apple iMovie. Unlike Experiment 1, the tone was only as long as the target word, not the entire utterance.

3.1.2.3.2. Linguistic Materials

As noted above, the linguistic materials differed from Experiment 1 in several ways. First, linguistic information was presented to the participants concurrently with the video in the form of closed captions. Automated captions were first generated using YouTube, and were then checked and corrected by human coders. Second, all speech within the 40 second video clip was captioned in Experiment 2, not only the utterance within which the target word appeared.

All words that appeared in the captions were tagged as either *content* or *function*. Five conditions of linguistic information were then created by manipulating which captions were presented.

- Full Video-and-Language (VL): closed captions for all speech other than the target word.
- Ordered Function Words (OFW): function words were captioned where they appeared in the utterance, while content words were removed. Each removed word was marked with “---”, and separated by spaces from other words.
- Ordered Content Words (OCW): content words were captioned where they appeared in the utterance, while function words and additional morphology (e.g., plural -s) were removed. Each removed word (but not the removed morphemes) was marked as in OFW.
- Alphabetized Content Words (ACW): the captions showed an alphabetized list of content words, separated by commas. Function words

and morphology were not included. A new list appeared for each utterance, as it was being said in the clip.

- Video-Only (VO): No captions appeared with the video.

These conditions allowed for the systematic manipulation of syntactic and content information, and resulted in the creation of 1615 unique versions of the clips [(3 Prohibition ‘Not’ instances * 5 Context Conditions) + (64 other Target Words and Negation Functions * 5 clips * 5 Context Conditions)].

In the Video-Only condition, in which no captions appeared, the target word was marked only by an auditory tone. In all other conditions, in addition to the same tone the target word was also replaced with ‘***’ in the captions. Each speaker was identified in the captions (e.g., “Mom:”). Any sung speech was surrounded by ‘♪’ symbols on both sides, while background music was not captioned. Table 2 displays an example of how a single utterance was captioned differently across conditions.

Linguistic Information Condition	Closed Caption
Original	Look at the houses and the trees and our car.
Full Video-and-Language (V+L)	Look at the houses and the trees and our ***.
Ordered Content Words (OCW)	Look --- --- house --- --- tree --- --- ***.
Ordered Function Words (OFW)	--- at the --- and the --- and our ***.
Alphabetized Content Words (ACW)	[House, Look, Tree]

Table 2. An example of how the same utterances were captioned in each linguistic information condition, excepting the Video-Only condition, in which no captions were shown. In all of the other conditions, captions appeared throughout the entire video, not just with the target word. The target word in this example was the noun “car”.

3.1.3. Procedure

The study was created and hosted on Qualtrics and published on Amazon Mechanical Turk. After participants consented to participate, they were asked to enter their age and whether they spoke any languages other than English. If they reported being under 18 or speaking any other languages, they were excluded from participation and directed to an exclusion debrief. Otherwise, participants saw an instruction screen. They were told that they would be watching videos filmed from a child’s perspective, and were told that they would be asked to guess what word was said when a beep appeared in each video.

To establish that they could hear the audio and see the video, participants watched a test video clip in which a sequence of 10 different words appeared on the screen, and a beep sounded over one of the words. They were asked whether they could hear the beep, view the video, and identify which word was printed during the beep. To progress to the study, participants had to enter the correct word and respond that they could hear and view the video.

Finally, participants viewed the 20 test videos, presented in random order. They saw one video at a time and could not repeat a video after it played. After each video, participants were prompted to enter their guess of the target word into a text box.

3.1.3.1. Randomization

Participants were randomly assigned to one of the five Language Information Conditions, one of the four Lists of words, and one of the five Clip Numbers nested within each List. This determined which target words they were tested on and which of five video clips they saw for each target word. Because there were only three clips of Prohibition *not*, these were randomly assigned to participants separately from Clip Number and List, so that each participant got one instance of Prohibition *not* (as well as each of the other combinations of negation word and negation function). The order in which video clips were presented was fully randomized between participants.

3.2. Results

All analyses followed a pre-registered plan (https://osf.io/r46jm/), except where otherwise noted. 312 individual trials (2.9% of 10,800 total trials across all participants) were excluded using criteria that were not pre-registered, but were considered before the data were plotted or analyzed: if the participant responded faster than 3 S.D. from the mean response time, gave multi-word responses, or responded inappropriately (with nonce words or expletives).

We used mixed effects logistic regressions to model the probability of guessing the target word correctly, and derived p-values for each of the effects that we report from Type II Wald chi-square tests comparing models with and without that effect. For each model described below, we started by including random intercepts for Subject wherever possible. Other Random intercepts (for Target Word or for Clip), as well as random slopes of each of these variables by each fixed effect and correlations between slopes and intercept were then added only where the more complex model converged and provided a significantly better fit based on model comparison (see Gries, 2015).

3.2.1. Effects of Information Type

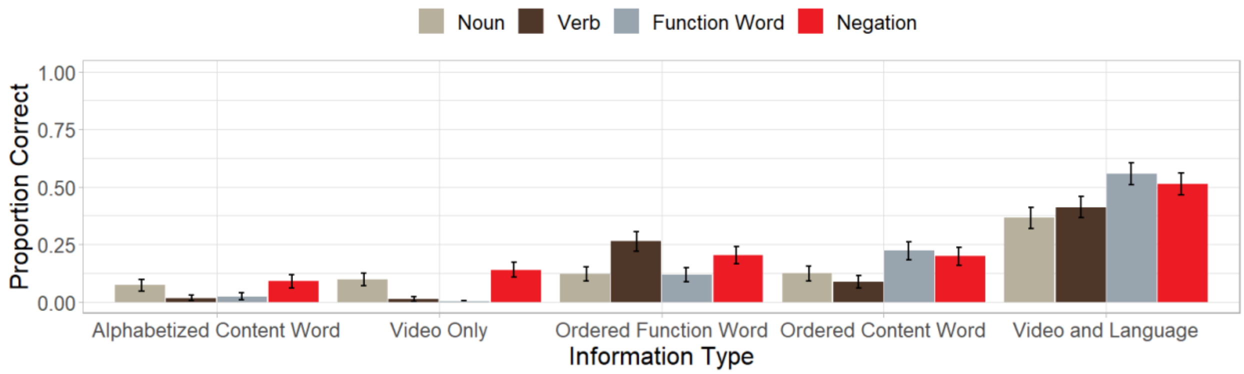

Figure 3: Proportion of correctly guessed target words, by Word Type and Information Type. All conditions included video. Error bars indicate ±1 Standard Error.

Figure 3 shows the proportion of trials in which participants correctly guessed the target word as a function of Information Condition (see Table 2) and Word Type (Noun, Verb, Function Word, or Negation). Although we had planned to code Information Type as a categorical variable, we found that analyzing the interaction of Information Type and Word Type was impossible under this coding scheme because even models with the most minimal random effect structure failed to converge. This led us to treat Information Type as an ordinal variable instead. We assigned each Information Type a numerical value based on the a priori expected relative difficulty of identifying the target Word. Given that we made no predictions about whether the target would be easier to identify in the Ordered Function Word or Ordered Content Word condition, we assigned them the same value [1 = Video-Only (VO); 2 = Alphabetized Content Word (ACW); 3.5 = Ordered Function Word (OFW); 3.5 = Ordered Content Word (OCW); 5 = Video-and-Language (V+L)]. The target Word Type was coded as a categorical variable with four levels. The only model that converged included a random intercept by subject.

This model revealed a significant main effect of Information Condition ($\chi^2(1) = 612.26$, $p < 0.001$), a significant main effect of Word Type ($\chi^2(3) = 70.78$, $p < 0.001$), and critically, a highly significant interaction between them ($\chi^2(3) = 136.45$, $p < 0.001$). To examine the main effect of Information Type and compare the observed effect to our predicted ordering of these conditions, we then ordered them by observed average success across Word Types. The observed ordering (V+L > OCW ⪆ OFW > VO ⪆ ACW) differed slightly from the one we predicted (V+L > OCW ≅ OFW > ACW > VO).

To compare participants' guessing success on each neighboring pair of Information Types in the observed ordering, we constructed four separate regressions, each restricted to subsets of the data corresponding to a pair of conditions, and each testing for a fixed effect of Information Type with a random intercept for the Subject and the Target Word. These analyses revealed no significant differences between participants' ability to guess the target word in the Alphabetized Content Word and Video Only conditions (β=-0.22, z=1.52, *p*=0.128, OR=0.8), or between the Ordered Function Word and Ordered Content Word conditions (β=0.14, z=1.36, *p*=0.175, OR=0.87). The main difference from our predictions was that being presented with an alphabetically ordered list of content words proved no more help to participants' guesses than seeing the video alone, suggesting that participants didn't benefit (and may have just ended up ignoring) an unhelpfully ordered list of content words presented at subtitle speed. While the Ordered Function Word and Ordered Content Word conditions were not significantly different from each other, participants were nearly four times more likely to guess the target word in the Ordered Function Word condition than in the Video-Only condition (β=1.35, z=11.31, *p*<0.001, OR=3.84), and six times more likely to guess the target correctly in the Video-and-Language condition than in the Ordered Content Word condition (β=1.8, z=16.01, *p*<0.001, OR=6.03). In sum, while more language information generally helped participants make correct guesses, clues about the semantic content of the utterance and about its syntactic structure were on average equally helpful.

Figure 4 shows participants' likelihood of guessing a negation word depending on whether a negation word had been the target, as a factor of the amount of information they received. To explore the interaction between Word Type and Information Type, simple effects analyses tested whether differences between Negation and other Word Types existed within each Information Type. Each analysis examined the subset of the data in each Information Type condition, coding negation as the reference level for comparison to each other Word Type. There were no significant effects of Word Type within the Video-Only ($\chi^2(3)$= 7.55, p = 0.056) or the Alphabetized Content Word conditions ($\chi^2(3)$ = 6.49, p = 0.090). In contrast, there was a significant effect of Word Type in the Ordered Content Word condition: ($\chi^2(3)$ = 11.92, p = 0.008), in the Ordered Function Word condition ($\chi^2(3)$ = 12.59, p = 0.006), and in the Video-and-Language condition ($\chi^2(3)$ = 10.31, p = 0.016), though in each of those conditions, participants' ability to guess negation words was not significantly different from any of the other word types. Given the difficulty of guessing "hard" words in the absence of linguistic context (cf. Gillette et al., 1999), it is striking that correctly guessing that the speaker had said *no* and *not* proved about as easy for participants as correctly guessing a common noun, verb, or other function word. Importantly, however, in grouping all types of negations together, these comparisons obscure what turn out to be large differences between different functions of negation.

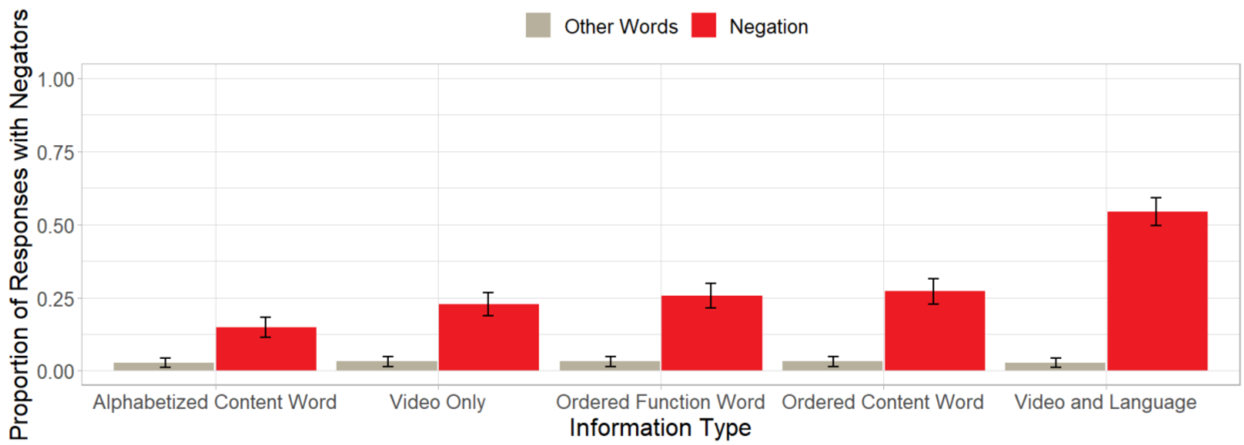

Figure 4: Proportion of guesses that were of any negation word (*no*, *not*, or any word ending in *n't*), as a function of whether the target word was actually a negation (*no* or *not*) versus any of the other word types (nouns, verbs, or function words), broken down by the amount of information. Error bars indicate ±1 Standard Error.

3.2.2. Analyses of Negation Word and Negation Function

The next set of analyses focused on examining differences between negation words (*no* and *not*) and between the communicative function these words were used to express (Denial, Prohibition, or Nonexistence), while disentangling these factors. Figure 5 shows how often participants were able to correctly guess each target Negation Word when it was used to express each Negation Function, divided across different conditions of linguistic information availability (i.e., Information Type). Each of the models below included a random intercept for Subject. Both because it was our pre-registered plan to do so and because no model with the three-way interaction of these factors converged, we analyzed each of the two-way interactions in separate models.

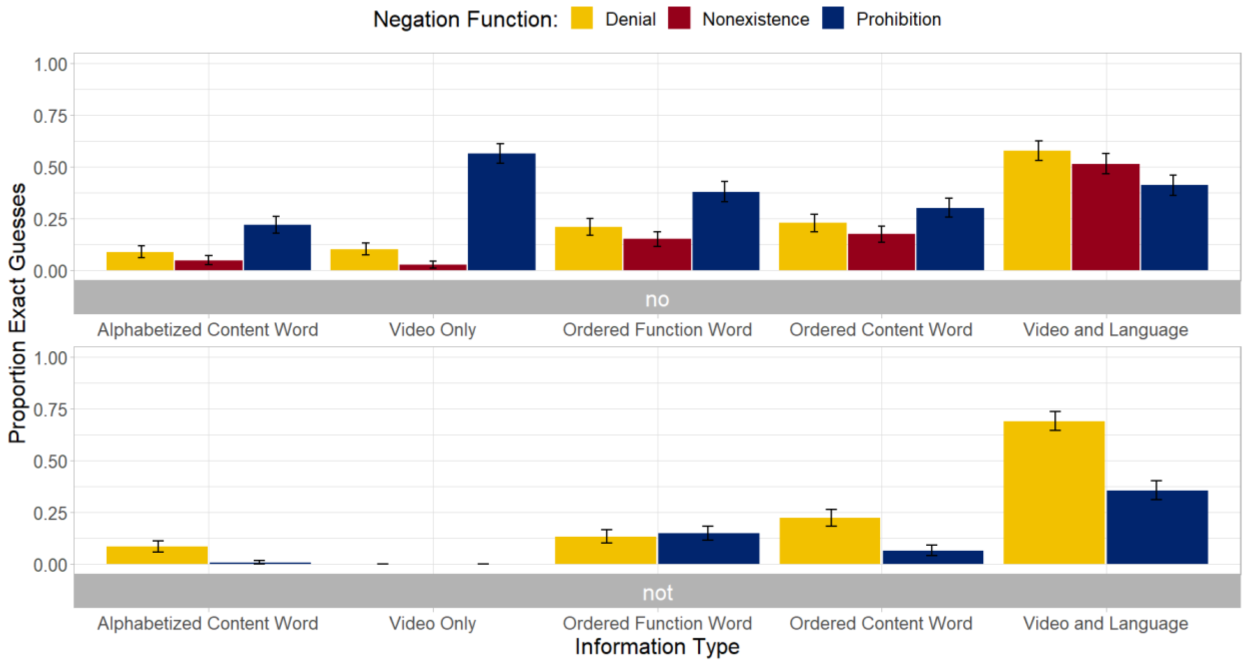

Figure 5: Proportion of correctly guessed target words, by Negation Word (above: *no*, below: *not*), the communicative function of that word (Negation Function: Denial, Nonexistence, or Prohibition), and the amount and type of language information provided (Information Type), along the x-axis, ordered by average number of correct guesses. All Information Type conditions included video. Error bars indicate ±1 Standard Error.

First, a mixed effects logistic regression probed the relation between Negation Word (*no* vs. *not*) and Information Type (modeled as an ordinal variable, same as above). This model found a significant interaction ($\chi^2(1)=60.16$, $p<0.001$), as well as significant main effects of Word ($\chi^2(1)=17.36$, $p<0.001$) and Information Type ($\chi^2(1)=141.35$, $p<0.001$), indicating that participants' odds of guessing *not* benefited more from additional language information than their odds of guessing *no*. Simple effects analyses further revealed that both words did benefit from additional language context, even if unequally (*not*: $\beta=1.25$, $z=10.49$, $p<0.001$, OR=3.48; *no*: $\beta=0.37$, $z=8.21$, $p<0.001$, OR=1.45).

Next, another model investigated the effects of Negation Function (Prohibition, Nonexistence, Denial; dummy-coded with Denial as the reference level), Negation Word (*no* and *not*), and their interaction. Recall that there were no instances of Nonexistence Not in the stimuli, so the interaction in the model is only of the contrast between Prohibition and Denial, interacting with whether the target Word was *no* or *not*. This model revealed a highly significant interaction between Negation Word (*no* vs. *not*) and Negation Function ($\chi^2(1)=50.53$, $p<0.001$), as well as significant main effects of Negation Word ($\chi^2(1)=48.58$, $p<0.001$), indicating that *no* was generally easier to guess correctly than *not*, and of Negation Function ($\chi^2(2)=32.93$, $p<0.001$), indicating that Prohibitions were generally the easiest to guess. Exploring the interaction further, simple effects analyses revealed that, when the target word was *no*, Prohibitions were significantly easier to identify than Denials, (β=0.69, z=4.79, p<0.001, OR=2.0) and Denials were in turn easier to identify than Nonexistence uses (β=0.40, z=2.51, p=0.01, OR=1.49). In contrast, when the target word was *not*, Denial uses were the ones that were easier to identify than Prohibitions (β=3.15, z=6.29, p<0.001, OR=23.38). The size of this effect was dramatic: participants were many times more likely to correctly guess a Denial use of *not* than a Prohibition use, considering only exactly correct guesses of the word.

However, further exploratory analyses revealed that this effect was driven by participants' tendency to guess *no* (472 guesses) rather than *not* (195 guesses) when the target was either of the two. If we instead credited participants for guessing any negation word when the target was any negator (e.g., guessing *no* or any word ending in *-n't* when the target word had actually been *not*), then differences between the functions of *no* persisted, but there was no longer any significant difference between Denial and Prohibition uses of *not* (β=0.13, z=0.88, p=0.379, OR=1.14). This suggests that the reason *not* was easier to correctly guess in Denial than Prohibition uses is that, when seeing video that looks like a parent is issuing a prohibition, participants tend to guess that the parent said *no* and hardly ever guess *not*.

Because children say *no* much earlier than *not* – starting at an age at which they are not yet able to make sense of much of the rest of their language input – we were particularly interested in investigating the interaction between the amount of language information and the function of negation (i.e. between Information Type and Negation Function) specifically when the target word was *no*, and specifically homing in on the identifiability of different functions of negation at low levels of language information. Looking at just the subset of trials on which *no* was the target word, we found that the main effects of Negation Function (χ^2(2)=35.14, p<0.001) and Information Type (χ^2(1)=37.79, p<0.001) were both significant, as was the interaction between them (χ^2(2)=88.06, p<0.001). Exploring this interaction further, we analyzed the effect of language Information separately within each Negation Function. We found significant effects of Information on participants' ability to identify Denial (β=0.75, z=7.94, p<0.001, OR=2.13) and Nonexistence (β=1.16, z=8.30, p=0.01, OR=3.20) uses of *no*, but importantly, no effect on their ability to identify Prohibition *no* (β=-0.09, z=-1.38, p=0.17, OR=1.10). Surprisingly, additional language context did not significantly help participants tell that a parent had said a Prohibition *no*. To further probe the difference between Prohibitions and other functions, we looked at the identifiability of the different functions of *no* within the two most maximally different Information conditions: Video-Only and Video-and-Language, with separate models for each of these conditions. We found that participants were significantly, many times better at identifying Prohibition than Denial uses of *no* in the Video-Only condition (β=2.48, z=6.51, p<0.001, OR=12), but significantly *worse* in the Video-and-Language condition (β=-0.68, z=-2.40, p=0.017, OR=1.97). We again see that for Prohibition, but not for Denial uses of *no*, visual information is all participants needed.

To investigate just how easy it was to identify Prohibition uses of *no* without any additional language information, an exploratory model compared participants' performance with Prohibition *no* (as the dummy-coded reference level) to their ability to identify each of the other word types (nouns, verbs, and function words) in the Video-Only condition. Each contrast in this model was significant, with large effect sizes. Without any language and with only visual information, participants were much more likely to identify Prohibition *no* than early common nouns (β=2.54, z=10.22, $p<0.001$, OR=12.68), many times more likely to identify it than early verbs (β=4.66, z=10.82, $p<0.001$, OR=105.29), and even more likely to identify Prohibition *no* than other function words (β=6.61, z=6.50, $p<0.001$, OR=745.35), which they virtually never guessed.

Finally, a third model probed the relation between Negation Function (Denial, Nonexistence, and Prohibition) and Information Type (modeled ordinally, as above). This model found a highly significant interaction ($\chi^2(2)=88.06$, $p<0.001$), and both significant main effects (Function: $\chi^2(2)=35.14$, $p<0.001$; Information: $\chi^2(1)=37.79$, $p<0.001$). Exploring the interaction further, simple effects analyses[4] found that each of the Negation Functions became easier to guess with additional language information (Denial: β=0.99, z=12.80, $p<0.001$, OR=2.69; Nonexistence: β=1.16, z=8.30, $p<0.001$, OR=3.19), though Prohibitions benefited the least (β=0.19, z=3.50, $p<0.001$, OR=1.21).

3.3. Discussion

The results of Experiment 2 replicate and extend the findings of Experiment 1. Distinguishing the function of negation from the word used to express it, we again find that Denial negations are especially difficult to identify in the absence of linguistic context and benefit the most from its inclusion, actually becoming the easiest use of negation to identify when participants are given full information about the surrounding discourse. Critically, Experiment 2 finds that these effects hold whether the word used to express Denial is *no* or *not*, showing that the pattern found in Experiment 1 was not just due to *not* being more frequently used to Deny and *no* being more frequently used to Prohibit. Rather, inferring that either *no* or *not* was used to Deny – that is, to express truth-functional negation – gets easier the more of the surrounding discourse a learner can understand.

In contrast to Denials, as well as expressions of Nonexistence, Prohibition negations were often identified correctly even in the total absence of any language information, and adding this information did not help to identify them. Whether *no* or *not*

[4] None of these models converged with a random intercept for Subject, probably because each Subject only contributed two data points for Denials and Prohibitions, and only one data point for Nonexistence 'Not.' However, all three models converged with a random intercept for the specific video clip that was used for each target word; this was the random effect in these three simple effect analyses.

had been used to prohibit, participants guessed that a single beeped-out word was *no* when seeing only an uncaptioned video of a Prohibition event. This suggests that Prohibition uses of *no*, in particular, have distinctive correlates with what the learner can perceive, and that learners can use this information to identify both the negator and its intended meaning even without understanding any of the rest of the language being spoken. Surprisingly, we find that in the absence of any language, Prohibition *no* was massively easier to identify than any other type of word we tested – by an order of magnitude easier than common early nouns, and by two orders of magnitude easier to identify than common verbs or function words. This could be why Prohibition *no* is one of the very first things most infants say – it is a usage they can understand when they cannot yet understand much else.

Finally, although the primary results of Experiment 2 match Experiment 1, the results of the two experiments did diverge in some respects. Statistically comparing responses across the Experiments isn't appropriate – participants in Experiment 1 heard a tone over the entire target utterance and, in the Video Only condition, were asked to guess the entire utterance; in contrast, in all conditions of Experiment 2, participants were asked to guess just one target word. Comparing relative differences across experiments is nevertheless meaningful. First, although Prohibition uses of *no* became easier to identify when linguistic information was added to video in Experiment 1, they actually became more difficult to identify with the addition of language in Experiment 2. Second, unlike in Experiment 1, Nonexistence uses of *no* in Experiment 2 patterned closely with Denial uses of *no*, being affected by the addition of language information in similar ways. Although it is impossible to tell which of the many differences between Experiments 1 and 2 might be responsible for these differences, one factor that might explain both of the diverging patterns is the difference in how much language information each experiment provided. In Experiment 1, the Video-and-Language condition provided participants with the sentence that included the target word, but no other context. In Experiment 2, the same condition provided participants with full captions of the entire 40-second video clip. If Prohibition *no* is easy enough to identify from the video alone – perhaps because its production is tightly time-locked to the child visibly doing something that invites prohibiting – then adding a large amount of language information before and after this event (as in Experiment 2) might be distracting instead of helpful. On the other hand, identifying a Nonexistence use of *no* might be helped by knowing more about the preceding discourse context. It might be easier to tell that a child was commenting on their juice running out if it is clear that the immediately preceding conversation had been about juice.

3. General Discussion

4.1. Summary: In support of Information Change

In two experiments, we find evidence that identifying the use of truth-functional negation (in Denial uses) proves challenging even for conceptually equipped learners when they are not given access to the surrounding linguistic context within which the negator occurred. Other, non-logical uses of *no* and *not* are easier. This pattern is predicted by Information Change accounts, but unexpected on any Conceptual Change account.

In particular, when the adults in our experiments simulated an early word learner, who has access to information about the visual world but barely any information about what anyone is saying, they rarely guessed that a negator had been said when it had been used to Deny. However, in this same simulation, adults were much more likely to guess correctly that the utterance contained a negator when it had been used to Prohibit. This finding suggests that a very young learner who only has access to the visual world could acquire the Prohibition function of the word *no*, which likely reflects a more limited, non-truth-functional meaning (see Feiman, et al., 2017). This could explain why children understand and produce only non-truth-functional uses of negation words as early as 12 months of age. It appears that Prohibition uses of *no* turn out to be strikingly easy to identify for those who understand no language at all.

In contrast, when adults simulated what it must be like for a more advanced learner, who has access to both the visual and linguistic context minus the negator itself, they were much better able to guess negators that had been used to Deny. This suggests that acquisition of the Denial function of negation words – the clearest (and maybe only) expression of the truth-functional operator – likely requires having already learned a great deal about other aspects of language, such as the meanings of other words and language-specific syntax. This could in turn explain why only older children (around their second birthday) understand truth-functional negation and use it to Deny.

Both experiments converged on this pattern despite a variety of methodological differences. Both experiments presented adults with videos in which negators were uttered by parents to their children in natural everyday circumstances in the home. These videos were muted, other than a beep placed during the target sentence, which participants were tasked with guessing. Beyond this similarity, the details of the designs of these experiments differed in many ways that might have mattered. Despite separating Negation Function and Negation Word (*no*, *not*) into two separate factors, using a first- rather than third-person point of view, providing participants with linguistic information from the entire vignette, placing a beep over the target word only, and asking participants to provide single-word responses, Experiment 2 found, just like Experiment 1, that adults are less able to identify Denial uses when they have less linguistic information. In contrast, both Experiments 1 and 2 found that adults who saw only an uncaptioned video were able to tell when a parent had used *no* to issue a prohibition. Indeed, in Experiment 2 Prohibition uses of *no* were the absolute easiest words to identify from video alone, easier even than common early-acquired nouns.

4.2. Could there still be conceptual change?

These findings indicate that there is no need to appeal to conceptual complexity in order to explain both why young children only use *no* to prohibit and why the truth-functional meaning of both *no* and *not* takes as long as another year to learn. Even conceptually sophisticated adults (who have been producing all attested uses of negation for at least 16+ years) nevertheless show the same trajectory when asked to learn these uses of negation again under simulated stages of information availability.

Nevertheless, it is important to highlight that although our results obviate the need to posit conceptual change, they do not rule out that it occurs, or that it could also be separately sufficient to explain the same trajectory. We did not 'simulate' in adults different abilities to represent the concept of truth-functional negation (nor can we imagine any way to do that), so we cannot determine if the same trajectory could be produced via simulated conceptual change. However, other evidence, from the study of homesign, suggests there may be contributions beyond Information Change to explain the acquisition profile of negation. In particular, Franklin, Giannakidou, and Goldin-Meadow (2011) examined in detail the production of negation from a child homesigner. They found a progression of negation use similar to that observed by Bloom (1968), with Rejection uses preceding Denials, even though the child in question was not learning someone else's language. This suggests that the Bloomian developmental pattern in production found in this child may not be due to the challenges of mapping form to meaning and hence to Information Change. On the other hand, information change is still a possibility, even in this case study. Of all the arbitrary signs this child might have used as a negator, what he produced was a headshake – a gesture also often used as a negator by adults. This raises the possibility that he learned this expression from the input provided by his parents after all, and incorporated the gesture into his homesign. In that case, it is also possible that, just like *no*, headshakes most clearly mean Prohibition or Rejection to a learner with no access to language input. Only later, once this child realized that truth-functional negation could be used for all of the same functions, the same headshake might gain the meaning of negation and be used for Denials. Both explanations of the acquisition trajectory of negation in homesign are live possibilities.

Nevertheless, our present findings do indicate that a child learner of English who can think negated thoughts would still face a tough challenge in having to identify which words express those thoughts, given initially limited language information. One possibility is that some form of conceptual change takes place earlier. If children cannot already think negated thoughts at birth, they must gain this ability by some means by the time they try to learn how those thoughts are expressed. In such cases, learning the word would be separate from, and conditioned on, forming the concept, which would need to be accomplished by some other process such as maturation, the development of auxiliary supporting capacities, or the language-independent assembly of the concept itself.

Another possibility, which relates word learning to concept learning, is that negation words are initially mapped onto some intermediary meanings that serve as precursors to truth-functional negation, and act as building blocks (or at least, as scaffolding) for the development of the full-fledged logical concept. One intermediary meaning could be the concept of contrariety, which does not follow the Law of the Excluded Middle as truth-functional negation does (see Feiman, Mody, & Carey, 2022). Another possibility is that the various early functions of negation – and their corresponding different meanings – are themselves the building blocks of the logical operator (for discussion, see McDermott-Hinman & Feiman, in press). Critically, any such account, which proposes some development from different conceptual primitives, must both explain how the parts can combine to make the full-fledged conceptual whole, as well as how experiences hearing and saying words would help. Future work is needed to see if such a developmental process can occur, let alone whether it actually *does* occur, in order to establish that conceptual change and word learning can work hand-in-hand in the emergence of a logical operator.

4.3. How is information used to infer negation?

Experiments 1 and 2 show that information about the visual context of an utterance and information about its surrounding linguistic context contribute differently to inferring that the utterance contains a negation word. The relative contribution of visual and linguistic information depends on whether that word expresses a truth-functional negation or a non-truth-functional use. But why does visual information help to identify non-truth-functional uses, and why does linguistic information help to identify truth-functional ones? Which aspects of a situation, and which aspects of language, drive learners to hypothesize that an utterance contains a certain use of a negator? Here we offer some proposals for each of the uses that these experiments examined.

4.3.1. Prohibition

In both experiments, we found that Prohibition was the easiest function of negation for participants to identify when given only visual information. Why is Prohibition – and especially Prohibition uses of *no* – so easy to identify? In addition to being expressible via negation, Prohibition is a speech act – an utterance intended to have illocutionary force – in which a speaker is attempting to get a listener to act (or stop acting) in a certain way (Austin, 1962). The perceptual correlates of Prohibition may therefore be clearer than those of Nonexistence or Denial. Visual information alone may be enough to tell children when they are doing something the parent does not want them to do, and when the parent seeks to intervene in some way. One potential source of information for Prohibitions might be gesture (e.g., head-shaking, finger-wagging). However, Experiment 1's corpus was coded for parental gesture and no relationship was found between Prohibition uses and increased use of gesture (see supplemental materials).

What other visual sources of information might help to infer that a parent is prohibiting their child? Our own informal inspection of the Prohibition videos in both experiments suggests that they show situations that involve a sense of urgency. They often involved swift movements to stop, remove and intervene, sometimes in a precarious or dangerous situation. It seems likely that observers are using these aspects of the visual context to infer the use of a negator, and that child learners might reasonably assume this as well. It's unclear whether it is the precarious situation, the parent's response, or both that might induce a Prohibition inference, since the two often co-occurred. For example, in one of Experiment 1's Prohibition vignettes, a child is standing on top of a dresser, and in another, trying to eat a Cheerio that fell on the floor. In both cases, the parent rushes to intervene. Similarly, in a Prohibition vignette in Experiment 2, the child reached for the father's phone; the father quickly took it away before the child could grab it and said "no" at the same time. However, even without the visual cues from the parents' action (rushing to pick them up; taking the Cheerio or the phone away), it might be sufficient for an observer (or the child themselves) to know whether the child is doing something that *should* be prohibited.

Strictly speaking, this finding suggests that a very young learner should acquire the Prohibition usage of *no*, rather than the Rejection usage. Though children under age two frequently use no to Reject and rarely to Prohibit, previous work suggests that the two functions are likely related in meaning. Introducing Prohibition into the taxonomy, Choi (1988) originally noted that Rejection and Prohibition largely differ in whether the negation is about something the speaker does not want to do or something they do not want their interlocutor to do. This raises the possibility that Rejection and Prohibition share a common meaning (perhaps negative affect, or a desire to make an event stop), such that – parents, beware! – understanding parents' Prohibitions might teach children to use *no* to Reject.

4.3.2. Denial

4.3.2.1. How Denial is not learned

In both experiments, we found that, in comparison to Prohibition, Denial uses of Negation were difficult to infer from visual context alone. It was only when the surrounding linguistic context was provided to the observer that Denial uses became easier to infer. This makes sense when we consider that unlike Prohibition uses of negation words, Denial uses are not consistently associated with any particular speech act and carry no particular illocutionary force. One can deny the truth of something without intending anyone to do (or not do) anything. Without the visual correlates of a particular speech act, what other cues could a learner use to identify the presence of a truth-functional negation?

One possibility is that the syntactic context of negators could aid in its acquisition. After all, it is well documented that the information provided by the syntax of an utterance constrains interpretations of more abstract words (e.g., verbs broadly, credal verbs in particular; Gleitman, 1990; Naigles, 1990; Yuan & Fisher, 2009; Gilette et al., 1999; Papafragou et al., 2007). Like the syntax of verbs, the syntax of English negation words is also highly constrained. Indeed, *no* and *not* both have uniquely identifying distributional profiles and syntactic properties. English *no* is one of only a few words that can appear as a single-word answer to a polar question, and the only one of those that is also a quantifying determiner (e.g. *no cats bark*). English *not* is an adverb that can only modify auxiliary verbs (as in ‘do-support’; e.g. *John does not run*) and, unlike the vast majority of other adverbs, cannot modify main verbs directly (**John not runs*). If learners could link these syntactic properties to the semantics of negation prior to learning the words themselves, the unique distributional profiles of *no* and *not* could help them infer that a negated meaning is being expressed by those words. Because these distributional properties are not shared by negators in other languages, infants might not be expected to possess such English-specific linking rules. Still, at least the English-speaking adults in our study should have clued in, if this kind of learning mechanism is ever available for inferring that negation is being expressed.

However, our results suggest that syntax on its own is not a reliable cue for inferring negation, even for adults. One specific aspect of the results of Experiment 1 illustrates this point. Recall that “matched affirmative” items were created as controls for negative sentences. For matched affirmative items, examples of parents producing affirmative sentences (“That is your ball”) were provided to participants with a blank inserted in a position that would allow for a negator (“That is ___ your ball”). Thus, rather than the blank replacing a word, a blank was introduced in the syntactic position where *not* can express negation in English. One might think that if syntactic evidence alone led to guessing that a negator had been used, then such affirmative controls would lead to negation guessing just as much as actual examples of negated sentences in which the negator had been replaced with a blank (e.g., “That is not him” -> “That is __ him”). Yet, participants were reliably more likely to infer the presence of a negator in actual negative sentences, as opposed to affirmative controls (Language Only, Denials: 69% negative guesses; Affirmative Controls for Denials: 11% negative guesses). If their syntax is the same, what other information could participants be using to distinguish Denials from Affirmative Controls?

4.3.2.2. How Denial is learned

We argued at the outset that learning to express negation requires reverse-engineering how a speaker is construing the world – a learner must realize that a speaker is expressing a thought in which negation is a constituent at the time when that speaker produces a verbal negator. One way learners might be able to reverse-engineer the speaker’s construal is by identifying contexts in which a speaker is especially likely to deny the truth of another proposition – what Wason (1965) first called “contexts of plausible denial”. As many researchers have found, these contexts arise when the preceding discourse raises a polar question (yes or no?) to which Denial is a plausible answer. Compared to a truth-functional negation being uttered out of the blue, negation in these contexts is easier to understand, is judged to be more felicitous, and is more likely to be used by a speaker in ordinary conversation (Dale & Duran, 2011; Nieuwland & Kuperberg, 2008; Orenes, Moxey, Scheepers, & Santamaria, 2016; Tian, Breheny, & Ferguson, 2010; Wason, 1965). Crucially, these effects have also been found in children as young as two years of age, who have only just begun producing Denials themselves (De Villiers & Flusberg, 1975; Hummer, Wimmer, & Antes, 1993, Nordmeyer & Frank, 2018b; Reuter, Feiman, & Snedeker, 2018), which suggests that the same contexts have the same effects on young learners as on adult speakers. For example, when shown five identical objects (e.g., five rubber ducks) and one different kind of object (e.g., a ball), both adults (Wason, 1961) and 2-year-olds (De Villiers & Flusberg, 1975) are more likely to say that the ball is “not a duck” than to describe one of the ducks as “not a ball”, or if they see just a ball without the ducks.

If children and adults both tend to deny in the same contexts, these contexts could form opportunities for children to learn which words are used to deny (that is, which words are used to express truth-functional negation). If a child who does not yet know how to express truth-functional negation can nevertheless think that the ball is not a duck, and if that child tends to think that thought in the same contexts in which adults tend to think it, and if, in that context, they hear an adult say, “the ball is not a duck”, word learning becomes an imminently solvable problem. The child does not even necessarily need to know what the adult is thinking, but could learn just by assuming that the adult is expressing the child’s own thought. With that assumption, the child only needs to determine which of the semantic constituents of their own thought maps onto which word the adult said. In shared contexts of plausible denial, the number of hypotheses a child needs to entertain about what *not* means shrinks from every concept they can entertain down to just the parts of the thought that they are having, and which they happen to be sharing with the adult speaker at that moment. Negators used in such contexts may then be “gems” (Gleitman & Trueswell, 2020), that is, negative utterances with high referential transparency that can serve as particularly excellent learning opportunities.

While the present evidence supports this possibility, it also highlights just how much it takes to identify the contexts of plausible denial in naturally occurring interactions between parents and children. In both experiments, we found that adult learners could not do it from visual information alone, and that they got better at identifying contexts of plausible denial the more they understood what else, other than the negator, was being said. This may likewise be the role that increasing language knowledge plays for child learners.

Visual information can also provide additional, sometimes unique clues that a context for plausible denial is present. In Experiment 1, we found that when the actual visual context was added to the syntactic frame of Denial utterances (comparing the Language-only to the Video-and-Language conditions), the likelihood of inferring that a negator had been used went *down* for Affirmative Controls but went further *up* for actual Negative sentences (Video-and-Language, Denials: 86%; Affirmative Controls for Denials: 7%). This suggests that visual information boosted accuracy in the identification of negation. In some rare cases, visual information might even be sufficient on its own. For example, Gomes et al. (2020) found that silent videos in which an action with a visually predictable end-state fails to achieve that state (e.g., a batter swinging at a ball and missing) regularly elicited negative descriptions, without any linguistic discourse context. The fact that failure is so salient in these events might set up a polar question under discussion (*did the event end as expected?*), much like a linguistic discourse would.

Still, such good visual correlates of Denials – ones that are sufficient to get the speaker and listener to think a similar negative thought at the same time – are rare, as evidenced by our finding that adults were very unlikely to infer a Denial given visual information alone. If early learning depends on perceptual observation, it is no surprise that a word that’s not usually amenable to learning by observation takes longer to learn. Comprehension of the surrounding verbal discourse is likely important to reaching a shared negative representation. It is possible that the same reasoning extends to other early logical words, with other “message-internal” meanings like conjunction or disjunction being difficult to acquire for the same reasons, and being acquired in similar ways. Note, however, that on this developmental account there is no point in development at which children’s conceptual system becomes capable of representing something it could not previously represent.

4.3.3. Nonexistence

As discussed in Section 3.3, participants' ability to infer the presence of negators used to express Nonexistence varied between the two experiments. As we suggested above, it is possible that Nonexistence uses have some key non-linguistic behavioral correlates (e.g., hand flipping, searching) as well as some interactions between language and scene (e.g., discussing something not present). Indeed, in our informal review of vignettes in both Experiments, this was often the case. However, performance in the Video-Only conditions of both experiments suggests that, even if present, the visual component of searching behaviors alone either does not appear distinctive enough, or is not specific enough to Nonexistence utterances. After all, there are many other utterances that might be just as relevant, or more helpful, for someone to say while searching (e.g., "where is it?"). Additionally, the relatively higher performance in Experiment 2 (where running captions were provided to participants) over Experiment 1 (where only the target utterance was provided to participants) in the Video-and-Language conditions, suggests that the absence of a referent on its own is not enough, and that discourse context may be most helpful for inferring talk about Nonexistence. This may be because it's actually quite often the case that non-present things are discussed, and so additional evidence beyond absence is needed to motivate the inference of Nonexistence negation. As we mentioned in the introduction, it is also possible that Nonexistence uses of *no* and *not* do not reflect a separate meaning from truth-functional negation, making them essentially unlike Prohibitions and more like a limited use-case of Denial. Whether Nonexistence uses express a separate meaning from truth-functional negation is an important topic for future work.

4.4. Conclusion

We asked what limits children's ability to learn that *no* and *not* express negation – their cognitive immaturity or their lack of language knowledge? We found that when adults – inarguably conceptually equipped learners – were restricted to watching a silent video of a parent and a child interacting, they struggled to infer that the parent had used a negator as a logical operator, to Deny. It was only with fairly complete linguistic information (every word but the negator) that they were able to reliably guess the presence of the negator. In contrast, when a negator had been used to Prohibit, we found that a silent video was as good as it gets – additional language information did not help participants guess that a parent had said *no* to Prohibit. These effects replicated across two independently conducted experiments, with many methodological differences between them. The robust pattern of findings suggests that changes in how much language children understand over time are enough to explain why it takes time to learn to express truth-functional negation. In other words, how children acquire the logical meanings of *no* and *not* might look very much the same if the concept of truth-functional negation were available to them from birth. These findings make it more plausible that children can represent truth-functional negation innately (though perhaps only implicitly, see Feiman et al., 2022 for discussion), bolstering other recent findings of children's early competence with negation (Cesana-Arlotti et al., 2018, 2020; de Carvalho et al., 2018; Feiman et al., 2022). It is a further open question whether children might have a similarly early competence with other logical operators.

Combined with similar findings that suggest Information Change may be the limiting factor for children's acquisition of other words (e.g., verbs generally (Gillette et al., 1999), and credal verbs in particular (Papafragou et al., 2007)), this study suggests a broader role for Information Change as a driver of vocabulary growth writ large. This in turn suggests that the concepts that different sorts of words express are available to learners very early in development.

Acknowledgements

This research was supported by a Fellowship from the Jacobs Foundation and a grant from the DARPA GAILA program, awarded to R.F., and by the National Institute of Child Health and Human Development Grant P01HD40605 (to S.G.-M.).

We thank the RAs that helped at various steps along the way: Sarah Fleming, Ebony Goldman, Jenna DiStefano at the University of Pennsylvania, and Allie Elkhadem, Haley Gibbs, Kristin Howell, Jennifer Hernandez-Pina, Naomi Lynch, and Max Mines at Brown University. Thanks also to Annika McDermott-Hinman for helpful comments and discussion. We would additionally like to thank Kristi Schonwald and the rest of the LDP team at UChicago for assistance in accessing and navigating the corpus in Experiment 1, and Asa, Erica Wojcik, and Jessica Sullivan for contributing to the SAYCam corpus, and for their generosity, advice, and assistance with respect to deploying it in Experiment 2. And, finally, we would like to thank Lila Gleitman, who helped us see the obstacles in learning by observation in the first place and provided insightful comments and incisive jokes throughout.

References

Ackrill, J., *Categories and De Interpretatione, translated with notes*, Oxford: Oxford University Press, 1975.

Austin, J.L., 1962. *How to do Things with Words*, 2nd edition, J.O. Urmson and M. Sbisá (eds.), Cambridge, MA: Harvard University Press.

Austin, K., Theakston, A., Lieven, E., & Tomasello, M. (20140609). Young children's understanding of denial. *Developmental Psychology*, *50*(8), 2061. https://doi.org/10.1037/a0037179

Bates D, Mächler M, Bolker B, Walker S (2015). "Fitting Linear Mixed-Effects Models Using lme4." *Journal of Statistical Software*, **67**(1), 1–48. doi:10.18637/jss.v067.i01.

Bloom, L. (1970). *Language development: Form and function in emerging grammars*. Cambridge, MA: MIT Press.

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

Caroll, Lewis (1871). *Through the Looking-Glass, and What Alice Found There*. London, UK: Macmillan.

Cartmill, E. A., Armstrong III, B. F., Gleitman, L. R., Goldin-Meadow, S., Medina, T. N., & Trueswell, J. C. (2013). Quality of early parent input predicts child vocabulary 3 years later. *Proceedings of the National Academy of Sciences*, *110*(28), 11278-11283.

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, 517–531. doi:10.1017/S030500090001254X

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

Dimroth, C. (2010). The Acquisition of Negation. In L. R. Horn (Ed.), *The Expression of Negation* (pp. 39–72). DE GRUYTER MOUTON. https://doi.org/10.1515/9783110219302.39

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

Fitch, A., Arunachalam, S., & Lieberman, A. M. (2021). Mapping Word to World in ASL: Evidence from a Human Simulation Paradigm. *Cognitive Science*, *45*(12). https://doi.org/10.1111/cogs.13061

Fox J, Weisberg S (2019). *An R Companion to Applied Regression*, Third edition. Sage, Thousand Oaks CA. https://socialsciences.mcmaster.ca/jfox/Books/Companion/.

Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2016). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*. doi: 10.1017/S0305000916000209.

Franklin, A., Giannakidou, A., & Goldin-Meadow, S. (2011). Negation, questions, and structure building in a homesign system. *Cognition*, *118*(3), 398-416.

Dale, R., & Duran, N. D. (2011). The cognitive dynamics of negated sentence verification. *Cognitive science, 35*(5), 983-996

De Villiers, J. G., & Flusberg, H. B. T. (1975). Some facts one simply cannot deny. *Journal of Child Language, 2*(2), 279-286.

Gleitman, L. (1990). The structural sources of verb meanings. *Language acquisition*, *1*(1), 3-55.

Gleitman, L. R., & Gleitman, H. (1992). A picture is worth a thousand words, but that's the problem: The role of syntax in vocabulary acquisition. *Current Directions in Psychological Science*, *1*(1), 31-35.

Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, *73*(2), 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. R., & Trueswell, J. C. (2020). Easy words: Reference resolution in a malevolent referent world. *Topics in cognitive science*, *12*(1), 22-47.

Goldin-Meadow, S., Levine, S. C., Hedges, L. V., Huttenlocher, J., Raudenbush, S. W., & Small, S. L. (2014). New evidence about language and cognitive development based on a longitudinal study: Hypotheses for intervention. *American Psychologist*, *69*(6), 588–599. https://doi.org/10.1037/a0036886

Gries, S. Th. (2015). The most under-used statistical method in corpus linguistics: Multi-level (and mixed-effects) models. *Corpora*, *10*(1), 95-125.

Hochmann, J.-R., & Toro, J. M. (2021). Negative mental representations in infancy. *Cognition*, *213*, 104599. https://doi.org/10.1016/j.cognition.2021.104599

Horn, L. R. (1989). *A natural history of negation*. Chicago, IL: 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

Landau, B., Gleitman, L. R., & Landau, B. (1985). *Language and experience: Evidence from the blind child* (Vol. 8). Harvard University Press.

McDermott-Hinman, A., Feiman, R. (in press). The development of negation in language and thought. Chapter in F. Blanchette & C. Lukyanenko (Eds.), *Perspectives on negation: Views from across the language sciences.* De Gruyter Mouton.

McNeill, D., & Mcneill, N. B. (1967). A question in semantic development--what does a child mean when he says" no". Studies in Language and Language Behavior,
Progress Report IV.

Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, *17*(2), 357–374. https://doi.org/10.1017/S0305000900013817

Nappa, R., Wessel, A., McEldoon, K. L., Gleitman, L. R., & Trueswell, J. C. (2009). Use of Speaker's Gaze and Syntax in Verb Learning. *Language Learning and Development*, *5*(4), 203–234. https://doi.org/10.1080/15475440903167528

Nieuwland, M. S., & Kuperberg, G. R. (2008). When the truth is not too hard to handle: An event-related potential study on the pragmatics of negation. *Psychological Science*, 19(12), 1213-1218.

Nordmeyer, A., & Frank, M. C. (2018a). Individual variation in children's early production of negation. In *CogSci*.

Nordmeyer, A. E., & Frank, M. C. (2018b). Early understanding of pragmatic principles in children's judgments of negative sentences. *Language Learning and Development*, 14(4), 262-278.

Orenes, I., Moxey, L., Scheepers, C., & Santamaría, C. (2016). Negation in context: Evidence from the visual world paradigm. *Quarterly Journal of Experimental Psychology*, 69(6), 1082-1092.

Papafragou, A., Cassidy, K., & Gleitman, L. (2007). When we think about thinking: The acquisition of belief verbs. *Cognition*, *105*(1), 125–165. https://doi.org/10.1016/j.cognition.2006.09.008

Pea, R. D. (1980a). The development of negation in early child language. In D. R. Olson (Ed.), *The social foundations of language and thought: Essays in honor of Jerome S. Bruner* (pp. 156–186). New York, NY: W. W. Norton.

Piccin, T. B., & Waxman, S. R. (2007). Why Nouns Trump Verbs in Word Learning: New Evidence from Children and Adults in the Human Simulation Paradigm. *Language Learning and Development*, *3*(4), 295–323. https://doi.org/10.1080/15475440701377535

Reuter, T., Feiman, R., & Snedeker, J. (2018). Getting to No: Pragmatic and Semantic Factors in Two- and Three-Year-Olds' Understanding of Negation. *Child Development*, *89*(4), e364–e381. https://doi.org/10.1111/cdev.12858

Russell, B. (1954). The metaphysician's nightmare. B. Russell, *Nightmares of eminent persons*, Allen & Unwin, London.

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

Sullivan, J., Mei, M., Perfors, A., Wojcik, E., & Frank, M. C. (2021). SAYCam: A Large, Longitudinal Audiovisual Dataset Recorded From the Infant's Perspective. *Open Mind*, *5*, 20–29. https://doi.org/10.1162/opmi_a_00039

Tian, Y., Breheny, R., & Ferguson, H. J. (2010). Why we simulate negated information: A dynamic pragmatic account.

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

Vilà-Giménez, I., Dowling, N., Demir-Lira, Ö. E., Prieto, P., & Goldin-Meadow, S. (2021). The predictive value of non-referential beat gestures: Early use in parent–child interactions predicts narrative abilities at 5 years of age. *Child development*, 92(6), 2335-2355.

Wason, P. C. (1965). The contexts of plausible denial. *Journal of verbal learning and verbal behavior*, 4(1), 7-11.

Yuan, S., & Fisher, C. (2009). “Really? She Blicked the Baby?”: Two-Year-Olds Learn Combinatorial Facts About Verbs by Listening. *Psychological Science*, *20*(5), 619–626. https://doi.org/10.1111/j.1467-9280.2009.02341.x

Yurovsky, D., Smith, L. B., & Yu, C. (2013). Statistical word learning at scale: The baby’s view is better. *Developmental Science*, n/a-n/a. https://doi.org/10.1111/desc.12036

Zehr, J., & Schwarz, F. (2022). *PennController for Internet Based Experiments (IBEX)*. https://doi.org/10.17605/OSF.IO/MD832

Zhang, Y., Yurovsky, D., & Yu, C. (2015, July). Statistical word learning is a continuous process: Evidence from the human simulation paradigm. In *CogSci... Annual Conference of the Cognitive Science Society. Cognitive Science Society (US). Conference* (Vol. 2015, p. 2793). NIH Public Access.

Supplemental Materials A

Coding Instructions

There are two negative function words in English:

• ***No*** •***Not*** (/ **-n't**)

But, we're focusing on 3 kinds of negation that have different communicative functions/effects:

- **Prohibition** - Stopping or altering child's behavior
- **Non-existence** - Noting the absence of something
- **Denial** - Asserting that something is not the case

No and *not* aren't cleanly divided into these kinds. Either can be used for any kind of negation listed above. Intuition says sentences with *no* as more likely to be prohibition or non-existence, and those with *not* as being denial, but we can't rely on this alone for categorization.

So, we'll be looking at three sources of information to help us determine which kind of negation is being used:

- **Linguistic** - the utterance itself
- **Immediate context** - discourse (e.g., surrounding sentences), co-present item/actions
- **Proceeding context** - how the parent or child responds to or reacts to the utterance

Linguistic information is the first thing you'll see in the spreadsheet, so from the initial sentence, you may be able to form a hypothesis about the kind of negation being used. But it may be ambiguous, e.g., "No candy," could be prohibition or non-existence. So, test your hypothesis by looking at the immediate and proceeding context. The preceding context may be useful as well, so feel free to check a minute or so before negation is uttered too.

Immediate context can tell you what is around (useful for non-existence) and what events are occuring (useful for prohibition) that might be relevant to the utterance. It can also provide additional communicative cues like prosody, joint attention, and gesture.

Proceeding context can help you discern the intended effects of an utterance. We typically say things for a reason, and how the infant, parent, or others in a video react to an utterance could tell us what they meant to do by saying something. For prohibition, there may be corrective, scolding, or redirecting behavior that follows from the parent, while there may be hesitation, stopping/changing tasks, or crying from the child. For non-existence, there may be seeking behavior or cleaning up. For both there may be crying.

On their own, each of the three can be ambiguous, but by **considering all three together, we can reduce the ambiguity**. There will most likely still be some extremely ambiguous utterances (e.g., "No," while the parent or child is out of view), so for those situations mark "unsure" and include a note.

Additionally, if the utterance seems to be saying something affirmative or rhertorical (e.g., "Aren't you adorable?"), mark "rhetorical" and include a note.

Immediate Context

Prohibition/Refusal	**Non-existence**	**Denial**
References things currently present/happening	Often references something not (/no longer) present	Correcting propositions/referring to untrue things
Attempt at stopping (picking up, taking something away)	Looking/seeking behavior	Can refer to something present, to correct
Finger wagging/head shaking	Hand flip gesture "I don't know/where is it gesture"	There may or may not be gesture
Sounding angry/irritated	Child upset	Indicating lack of knowledge or ability

Proceeding context

Prohibition/Refusal	**Non-existence**	**Denial**
Parent taking thing away or stopping action	Parent giving child more of item	Showing/getting /doing correct thing
Parental scolding	Taking away container for cleaning/disposal	Putting X back ("That's not your X")
Child stops/corrects action	Followed by asking "where"	Saying affirmative version ("This is X")
Child upset	Looking/seeking/replacing	Confusion or correction

Experiment Instructions

Video-Only Instructions

Your task in this study is to guess what words a parent has said in a muted video. We are interested in how children learn their first words, by noticing the situations in which the words are uttered by adults. To help us investigate this, you will be watching a series of short videos in which parents interact with their young children, speaking naturally and using the kinds of sentences that are typical when talking to small children – simple sentences with simple, common words.

In each video, we have picked out one sentence that occurred in the parents' speech to their child that we want you to try to guess. The sound will be absent from these videos, so you will not be able to hear what the parent was actually saying. However, at exactly the moment when the parent uttered the sentence of interest, we inserted an audible beep. Your task is to guess what the parent said during that beep.

After each video, you will see a text box to type in your guess, for example, if you thought the parents said “The doggie is sleeping,” you would type “The doggie is sleeping” into the text box.

Please do not leave any blanks – you should make a guess after each video.

To summarize, you will see a series of short videos (27 videos total) of parents interacting with their children in their homes. These videos will be muted, and you will hear a beep that tells you when the parent said the sentence you are supposed to guess. Immediately after the video, you will be given a text box to type in your response. Once you submit your answer, the next video will begin. Videos will only play once.

Make sure you are ready to start the experiment. When you are, click "Begin" below and it will start. Pay attention because the videos are short and only play ONCE.
Language-Only Instructions

Language-Only Instructions

Your task in this study is to guess, based on a partially complete sentence, what word a parent used in interacting with their child. We are interested in how children learn their first words, by noticing the situations in which the words are uttered by adults. To help us investigate this, you will see a series of sentences taken from videos in which parents interact with their young children, speaking naturally and using the kinds of sentences that are typical when talking to small children – simple sentences with simple, common words.

In each sentence, we have picked out one word that occurred in the parents' speech to their child that we want you to try to guess.

For example, if the parent had said "The nice doggie is sleeping," you would see a hint like “The ___ doggie is sleeping,” to help you guess what they said. Or, if they had said "I'm gonna go now," you would see a hint like "___ gonna go now." Your task would be to fill in the blank with a word in English that you think could go there. Sometimes the parent says just a one word sentence. When this happens, your clue will look like this: The parent said "_____." In this case, you should just guess the single word that the parent said.

"Please do not leave any blanks – you should make a guess for each sentence.

To summarize, you will see a series of sentences (27 sentences total) taken from videos of parents interacting with their children in their homes. The sentence will have a blank where the missing word should go, and you will be given a text box to type in your response. Once you submit your answer, the next sentence will appear.

Please do your best to guess what the parent said to their child.

Video-and-Language Instructions

Your task in this study is to guess what words a parent has said in a muted video. We are interested in how children learn their first words, by noticing the situations in which the words are uttered by adults. To help us investigate this, you will be watching a series of short videos in which parents interact with their young children, speaking naturally and using the kinds of sentences that are typical when talking to small children – simple sentences with simple, common words.

In each video, we have picked out one sentence that occurred in the parents' speech to their child that we want you to try to guess. The sound will be absent from these videos, so you will not be able to hear what the parent was actually saying. However, at exactly the moment when the parent uttered the sentence of interest, we inserted an audible beep. Your task is to guess what the parent said during that beep.

After watching each video, you will be given a hint as to what the parent actually said. We have removed ONE WORD from the parent's sentence. For example, if the parent had said "The nice doggie is sleeping," you would see a hint like "The ___ doggie is sleeping," to help you guess what they said during the beep. Or, if they had said "I'm gonna go now," you would see a hint like "___ gonna go now." Your task would be to fill in the blank with a word in English that you think could go there. Sometimes the parent says just a one word sentence. When this happens, your clue will look like this: The parent said "_____." In this case, you should just guess the single word that the parent said.

Please do not leave any blanks – you should make a guess after each video.

To summarize, you will see a series of short videos (27 videos total) of parents interacting with their children in their homes. These videos will be muted, and you will hear a beep that tells you when the parent said the sentence you are supposed to guess. Immediately after the video, your hint will appear under it, and you will be given a text box to type in your response. Once you submit your answer, the next video will begin. Videos will only play once.

Make sure you are ready to start the experiment. When you are, click "Begin" below and it will start. Pay attention because the videos are short and only play ONCE.

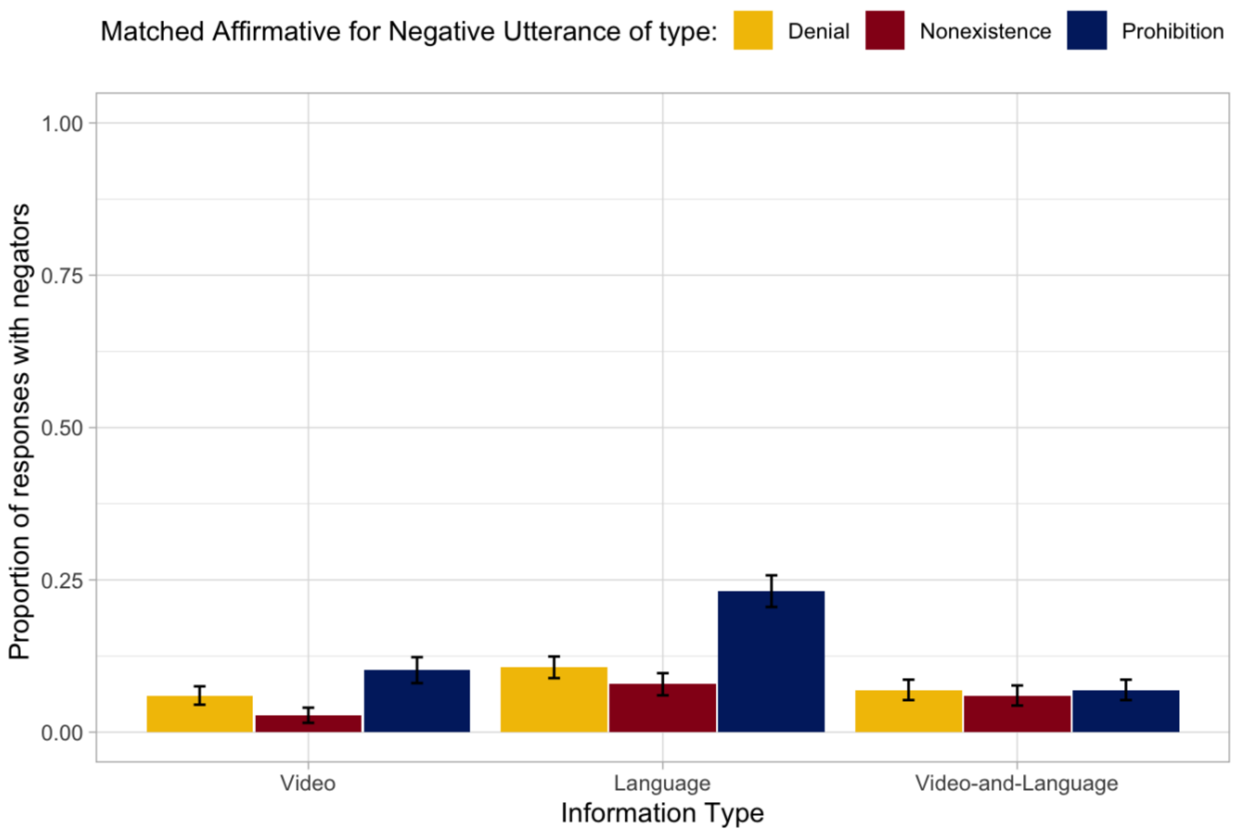

Figure S1: Proportion of negators included in participant responses for matched affirmative vignettes by the Negation Function that each affirmative is paired to. Data are presented by Negation Function (Denial, Nonexistence, Prohibition) and Information Type (Video Only, Video + Language, Language Only), with error bars indicating ±1 Standard Error. As described in Experiment 1: Methods, for each negative utterance, a non-negative utterance was sampled based on the presence of the word negated in the original negative utterance. This procedure did not take into account the "type" of affirmative utterance (so, e.g., a matched affirmative for a Nonexistence need not to assert existence explicitly).

Supplemental Materials B

The corpus used in Experiment 1 (the LDP corpus) was pre-coded for gesture (Goldin-Meadow et al., 2014) which allowed us to quantify the gestures used by the parents in our subsample of the corpus as described in section 2.1.2.2. The gestures of interest were determined based on those analyzed in Franklin et al., 2017. Namely, these were shaking of the head and a flip of the palm upwards. These gestures were fairly rare, only occurring 43 times throughout the entire transcript. Out of 9,841 utterances total, 789 included a negator (no, not, or -n't) and 24 of these utterances were accompanied by a gesture. Table S1 reports the counts of each gesture by Negation Function with the final row reporting gesture counts for utterances that did not include any negator. A post-hoc analysis using a mixed effects model to predict the presence of a negative gesture by whether the accompanying utterance included a negator found that parents were significantly more likely to use a negative gesture if the utterance included a negator than if it did not (β=1.33, z=8.67, $p<0.001$).

Function	Shake	Flip	Total
Prohibition	10	0	10 (out of 354)
Nonexistence	1	1	3 (out of 59)
Denial	10	1	11 (out of 375)
Affirmatives	2	17	19 (out of 9089)

Table S1: Counts of gestures that occurred with an utterance that included a negator, and all affirmative utterances. The values in parenthesis in the final column indicates the total number of utterances for each type.