HOW TO LEARN *NOTHING*: NEGATOR LEARNING AS A MAPPING PROBLEM

Victor Almeida Rodrigues Gomes

A DISSERTATION

in

Psychology

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2024

Supervisor of Dissertation

John Trueswell

Professor in Psychology

Graduate Group Chairperson

Russell Epstein

Professor in Psychology

Dissertation Committee

Anna Papafragou, Professor in Linguistics

Charles Yang, Professor in Linguistics

HOW TO LEARN *NOTHING*: NEGATOR LEARNING AS A MAPPING PROBLEM

Dedicated to those that love to play and only do their best when they are told they oughtn't.

ACKNOWLEDGMENT

If it's true that each something is a celebration of the nothing that supports it, then this work proves the converse is true, as it is certainly supported by many somethings. First and foremost, I would like to thank my mother, without whom I would not. She has a habit of making something out of nothing in spite of what logic, or anyone, has to say about it. I would like to thank my friends who helped me throughout this experience: Julian Shendelman, for being perhaps the best man I've ever met; Zach Ozma, for always reminding me to not overintellectualize this experience; Ally Paul, who I sometimes feel has done more for my future than I have and yet still seems to enjoy spending time with me; Gizmo, who I'm convinced knows what *no* means but simply won't take it for answer. I would like to thank my former colleagues at the Mütter Museum's Center for Education for reminding me about what really matters, and all my former out4STEM students for giving me hope. Relatedly, I'd like to thank the members of the Neo-Luddite Reading Group (and Iffy Books, our host) for being down to talk about seemingly hopeless issues. Having some company at least made it more bearable, in my opinion.

I would like to thank all of my teachers that believed in me, even though that usually just translated into more work. However, I would like to thank, Roger Reetz, my 7th grade math teacher who affectionately referring to me as a "shaggy headed idiot," for not believing in me. I doubt I would have gotten this far without someone to prove wrong. I'd like to thank Vovó Dalva, my first teacher to become like family, for showing me all that education could be. I'd like to thank my professors, especially Allen Schneider, for taking a chance on a freshman who showed up "too early in a lot of ways"; Dan Grodner, for formally (well, pragmatically) introducing me to psycholinguistics; Peter Bauman for his general commitment to curiosity and good conversation. Despite not knowing him personally, I'd like to extend a thanks to Professor Wala Kitu for taking *nothing* seriously. I'd like to thank Anna Papafragou for being on my committee (yet again!), and Charles Yang, for serving on my committee and keeping me (computationally) honest. I'd also like to thank Lila Gleitman: thank you for making me feel like a part of an age-old conversation, and somehow tricking me into actually speaking up and saying something. And, of course, I'd like to give a special thanks to John Trueswell for putting up with my frequent bouts of negativity. Of course, I know he didn't see it that way, and really he did so much more: he took me seriously and was never dismissive of what I cared about or what I was going through, but at the same time encouraged me to slow down and give everything it's time. But, listing all of his affirmative contributions would both take too long in addition to being significantly harder to pun around. My labmates, many of whom I also shared an office with: Caroline Beech, Spencer Caplan, Alex de Carvalho, Monica Do, Alon Hafri, Martin Ip, Tyler Knowlton, Sandy LaTourette, Abigail Laver, Zoe Ovans, and Heesu Sun. I'd like to thank all the brilliant undergrads I worked with, either as mentees or simply collaborators, especially: Jenna DiStefano, Allie Fehr, Sarah Fleming, Ebony Goldman, Yubin Huh, Abigail Laver, Marianna Olave, Arina Paniukhina, Rudmila Rashid, and Katinka Tangen. And, of course, I must thank my absolutely fantastic cohortmates, Max Genecov and Samiha Islam, for making my experience outside of the lab much funnier and on the whole more tolerable. I would like to extend deep gratitude to Kathrine Timmons, Heather Calvert, Jessica Marcus, Cynthia McElmore and Mark Liberman for regularly swooping in and offering some food or having a meal with me. I don't know why they did that, I suppose it seemed small at the time, but that couldn't have been further from the truth.

Finally, I'd like to thank everyone who told me no, be it the participants in my study or the administrators leafing through my applications. Without them, the yeses simply wouldn't have been much help. To any future scholars of cognitive science who similarly decide to ignore Parmenides' advice: if anyone accuses you of having a negative project, calmly remind them that the urge to negate is also a creative urge. ;)

ABSTRACT

HOW TO LEARN *NOTHING*: NEGATOR LEARNING AS A MAPPING PROBLEM

Victor Gomes

John Trueswell

Given its complexity, abstractness, and central role in many logics, negation might be a conceptual accomplishment. Therefore, young children's gradual acquisition of negation words might be due to their undergoing a conceptual change that is necessary to represent logical meanings. However, it's also possible that expressing negation takes time because children are gradually getting a 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 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 parents using negators when speaking to their children, adults inferred when the parent was used negation to prohibit but struggled with logical negation. However, when provided with additional linguistic information, guessing that the parent had expressed logical negation became easy. Thus, previously noted patterns in negator learning can be explained as a result of information access (i.e., learning more about one's language). Despite the effect of language, the experiment suggests there is still a contribution of the context. As such, it becomes critical to determine what sorts of contexts are reliably likely to encourage "negative thoughts." Eliciting negators has been challenging and has often been assumed to require explicit contrast with affirmative alternatives. This would suggest that eliciting negators may require explicit contrast with an affirmative linguistic variant and therefore more contextual support. Yet, adults are known to naturally produce negation in null-discourse contexts ("It isn't raining!"), and learners obtain command of sentential negation fairly early (possibly by 18 months). Inspired by work on event perception, we explore the possibility that prediction failure in event perception is such a null-discourse context and provide evidence that this holds across development. Finally, connections are made to pre-linguistic research and it is argued that logical words face the same issues as non-logical words, and therefore nativist positions are warranted on the same grounds.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

"Defenceless under the night
Our world in stupor lies;
Yet, dotted everywhere,
Ironic points of light
Flash out wherever the Just
Exchange their messages:
May I, composed like them
Of Eros and of dust,
Beleaguered by the same
Negation and despair,
Show an affirming flame."
- W.H. Auden, 1940

Much of the literature on concepts (i.e., how word meanings are represented) has focused on those concepts that are thought to pick out categories of things in the world. Though the exact phrasing and implementation can vary (concepts are "mental categories" (Gleitman et al., 1983), are "of mental categories" (Murphy, 2010) or that concepts include categories (Malt, 1990)), there is general agreement that concepts are implicated in categorization. The focus on such words is likely in large part because they seem more amenable to observation and experimentation (e.g., sorting task) for both conceptual development and word learning researchers. Unfortunately, this focus has left the field ill-equipped to discuss words that do not seem to correspond to categorical concepts, like *not*. This question is especially pressing for understanding the development of mathematical (see Carey, 2009 for review) and logical concepts (Bermúdez, 2003; Mody & Carey, 2016; Feiman et al., 2017; Cesana-Arlotti et al., 2018; Feiman et al., 2023). This issue has received increasing attention, likely in part due to the difficulties faced in uncovering the conceptual structure for words with seemingly simple categorical meanings like *dog* (Fodor, 1975). The definition of formal concepts, like negation or number, is the object of study of mathematics and logic, however, and so there are a range of candidate concepts more amenable to such study. This dissertation will argue that the same issues arise in uncovering the conceptual structure of formal concepts. In particular, a chief failing of earlier attempts to conclusively determine conceptual structure is that it is difficult to tie to a theory of word learning. This was noted in response to one of the earliest such theories of conceptual structure by Locke (1690), where Berkeley (1710) notes:

> "…the natural Consequence should seem to be, that so difficult a thing as the forming abstract Ideas was not necessary for Communication, which is so easy and familiar to all sorts of Men. But we are told, if they seem obvious and easy to grown Men, It is only because by constant and familiar use they are made so. Now I would fain know at what time it is, Men are imployed in surmounting that Difficulty, and furnishing themselves with those necessary helps for Discourse. It cannot be when they are grown up, for then it seems they are not conscious of any such Pains-taking; it remains therefore to be the business of their Childhood. And surely, the great and multiplied Labour of framing abstract Notions, will be found a hard Task for that tender Age. Is it not a hard thing to imagine, that a couple of Children cannot prate together, of their Sugar-plumbs and Rattles and the rest of their little Trinkets, till they have first tacked together numberless Inconsistencies, and so framed in their Minds abstract general Ideas, and annexed them to every common Name they make use of?"

Similarly, recent developmental psycholinguistic research suggests that some logical concepts, like negation, are learned far earlier than previously assumed (de Carvalho et al., 2018). Such work suggests that despite their complexity, negators (e.g., *not* and *no* in English) are learned alongside other early words, making it, like "Sugar-plumbs and Rattles," unclear when children have the time to learn such a potentially sophisticated concept. Indeed, achieving an active understanding of negation is a spiritual accomplishment in some belief systems, while in others it remains at least an academic accomplishment even if not a real one. On the other hand, having the ability to negate seems almost guaranteed by the basic fact that human cognition is fallible (Bacon, 1620). If anything is innate, after all, certainly it's the ability to be wrong. This dissertation will attempt to argue against conceptual stage views of negation by developing a clearer notion of the problems negator learning would present to pure "mapping" approaches (i.e., all that has to be learned is which label goes with which concept). Chapters 1 and 2 will present experimental data focusing on the contribution of language and observation to negator learning respectively, while Chapter 3 will return to basic questions about the relationship between conceptual structure and discontinuity. Chapter 3 aims to synthesize a nativist position on negation in light of these continuing difficulties.

Chapter 1 provides experimental evidence that previously documented patterns in the production (and comprehension) of negation are not prima facie proof of conceptual discontinuity. Early research in developmental psycholinguistics found that early uses of negators are typically restricted to simpler and more specific communicative functions (often Rejection ("No veggies!"), or sometimes Nonexistence ("No doggy.")), with uses more comparable to truth-functional negation (i.e., Denials) only occurring much later (~18 months vs. 2.5 years; Bloom, 1970; Choi, 1988), and later research suggesting Denial comprehension similarly arrives later (around 2 years; Feiman et al., 2017; Austin et al., 2014). As a result, some have argued that this discontinuity in linguistic behavior may be a reflection of an underlying discontinuity in conceptual competency (Dimroth, 2010; Pea, 1980; Hummer et al., 1993). Since the early vocabulary is dominated by nouns, similar arguments have been advanced in the past for the relative complexity of verbs meanings over noun meanings (Gentner, 1978; 1981; Carey & Levine, 1982). Others argued in response that the noun-dominant vocabulary is a consequence of early word learning depends mostly on observing co-occurrences between labels and referents (Gillette, 1999; see Gleitman et al., 2005 for a review), and it's only by learning more about their language (e.g., syntax) that they can take advantage of existing linguistic evidence to motivate word learning hypotheses. To test the possibility, Gillette et al. (1999) developed the Human Simulation Paradigm, where adult participants "simulate" learners who were artificially restricted to particular subsets of information present in word learning contexts. The participants were shown short (40s) videos of parents uttering some early word (a noun or a verb); these videos were muted other than a beep at the target word. Because the adult participants can inarguably think about events, they can simulate hypothetical learners that are limited only in their ability to access the linguistic information present in the scene. Despite their conceptual competence, adults still struggled in learning verbs from naturalistic verb contexts unless they were provided with additional syntactic information. The experiments in Chapter 1 extend this paradigm to investigate the possibility that Denials are seemingly delayed in production in comprehension because they require access to more sources of linguistic information to infer their uses than uses of other kinds of negation (i.e., Prohibition, Nonexistence). These predictions were confirmed, suggesting that conceptual change is similarly not necessary to account for the seeming delays in negator production and comprehension. Instead, it is likely that children are solely limited by the sources of evidence they can access as they learn a language (e.g., syntax, morphology, etc.).

Though gaining access to linguistic information is sufficient to explain the discontinuity in linguistic behavior, some of the data reviewed in Chapter 1 support the intuition that linguistic information alone is not sufficient to make negation relevant to a word learner. Because infelicitous uses of negation are difficult for children to process (Reuter et al., 2018; Nordmeyer & Frank, 2015; 2023), and the felicity conditions of negation can at times be complex, it's critical to investigate the sorts of null-discourse environments that might regularly occur throughout the process of word learning. Chapter 2 draws influence from the event perception literature and argues that failures may be such ideal "gems" (Gleitman & Trueswell, 2020) for mapping *not* onto negation initially. In two production experiments, both children (4-7yo) and adults were significantly more likely to negate when describing an event fail than when describing it succeed - and this result is not dependent in any way on local contrast (i.e., seeing the success before the failure). Though additional work is required to confirm that such a relationship to event perception would be of use to the youngest learners, Chapter 2 presents compelling evidence that failures lead both young children and adults to think of and produce negation in a strong and reliable manner.

Chapter 3 returns to the fundamental issues introduced in the beginning of this introduction and attempts to develop and extend the notion of access hinted at in Chapter 1. As such, Chapter 3 spends substantial time reviewing both the previous conceptual literature and recent non-linguistic research on negation. It then focuses on a particular, well-refined, proposal for the definition of negation built out of contrariety (McDermott-Hinman & Feiman, 2023) and argues that it does not overcome the known issues to definition. This is because the definition remains complex, and because logical words do not refer, it is unclear what sorts of environments or tasks not only can be solved by negation, but would be obviously solved by negation to a very early learner. That is, it requires a theory of access for it to figure in word learning, much like the definitions of categorical concepts, and this remains lacking despite the formal definability of logical words. Relatedly, the non-linguistic experiments can only speak to conceptual competence if one assumes that the tasks are not only solvable with negation (and not e.g., object tracking) but to some extent obviously solvable with negation. That is, that someone with the ability to represent negation would be compelled to do so in the task in question. However, it’s very likely that the contexts in which negation is accessible to infants change over time. As Chapter 1 suggests this could be because of linguistic scaffolding, but Chapter 3 will propose that with a label in hand, simple association could make the learner aware of contexts in which it’s useful to access a concept, and therefore influence non-linguistic behavior.

CHAPTER 1: IT'S NOT JUST WHAT WE DON'T KNOW

THE MAPPING PROBLEM IN THE ACQUISITION OF NEGATION

This chapter was previously published in 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. Experiment 2 was conducted and analyzed by collaborators at Brown University, however the rest (including the write-up of both experiments) was conducted by the author of the dissertation.

"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."
- L. 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 is not 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.

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, 2017) 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 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-

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 cannot 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).

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 et al., 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, 2018) 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 have not 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.

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.

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 wouldd 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.

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

[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.

alone would be sufficient to explain why only older children produce Denials and understand *no* and *not* as expressing truth-functional negation.

Experiment 1

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

Methods

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 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.

Materials

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.

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 would 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, 2018), 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.

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 Appendix A). 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%.

Stimulus Construction

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).

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 ___”).

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 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	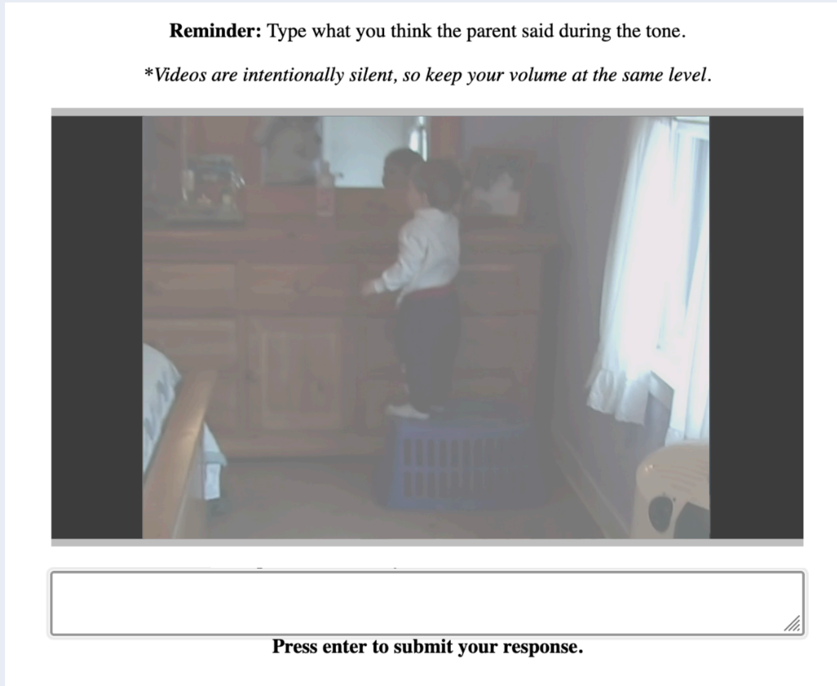
Video-and-Language Condition	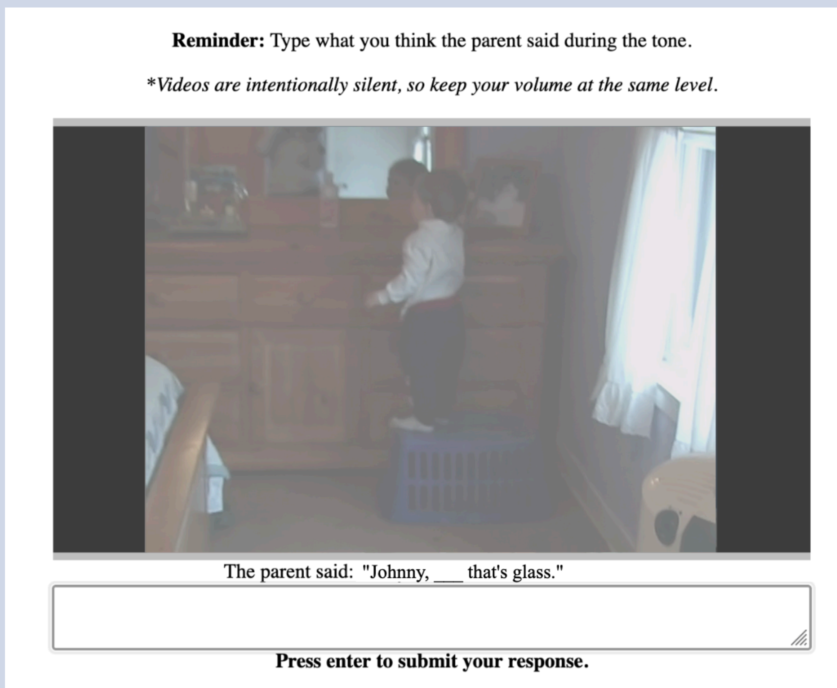

Language condition	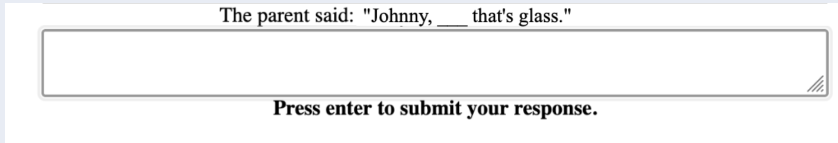

Table 1: The final screen of a trial shown in each Information condition. For the conditions that included Video, the textbox (and the linguistic stimuli in the Video & Language condition) only appeared after the video finished playing. After the video played once, it would not play again.

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.

Analysis

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.

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.

Results

Negation overall

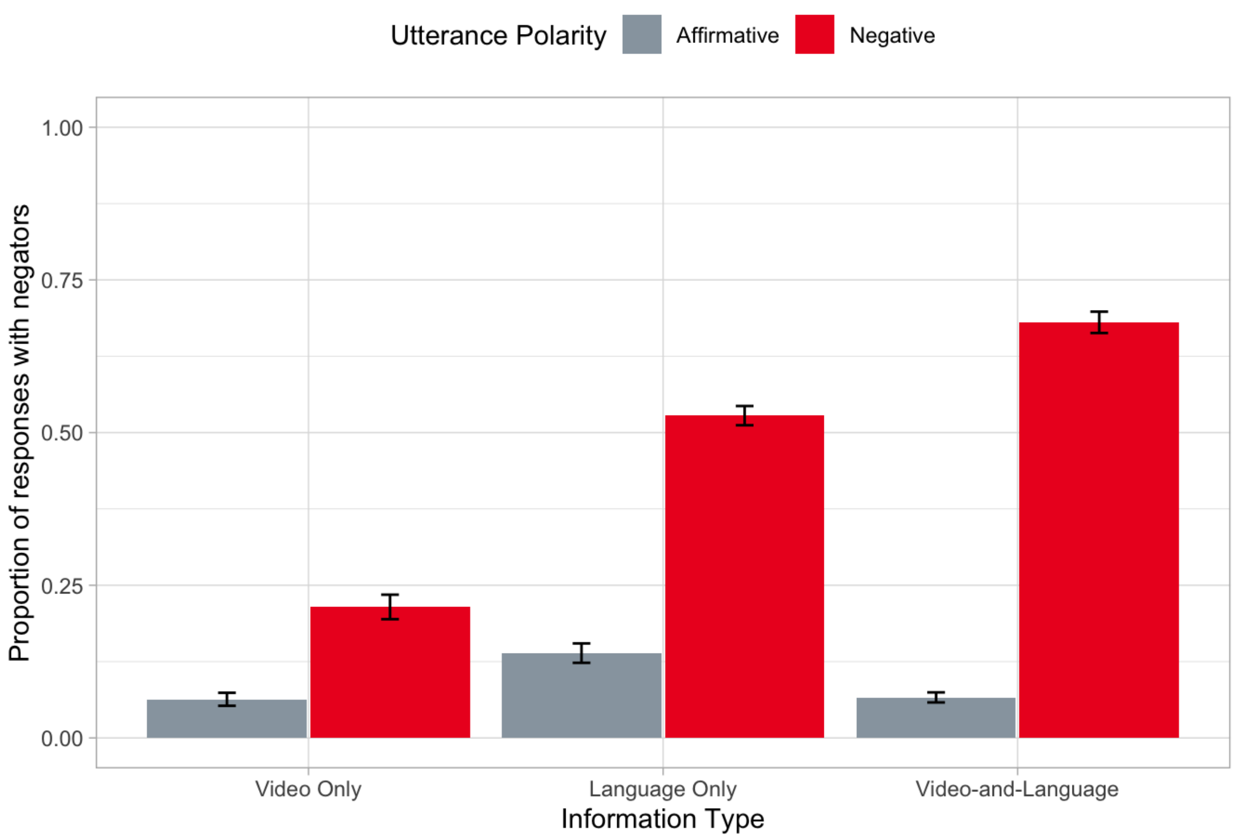

Figure 1: Proportion of negators included in participant responses for all Information Type conditions. Data are presented by Utterance Polarity (Negative, Matched Affirmative) and the amount of information provided along the x-axis (Information Type) along the x-axis. Error bars indicate ±1 Standard 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, χ^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 ($\beta=0.55$, $z=8.36$, $p<0.001$, 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 ($\beta=1.41$, $z=4.14$, $p<0.001$, OR=4.08).

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).

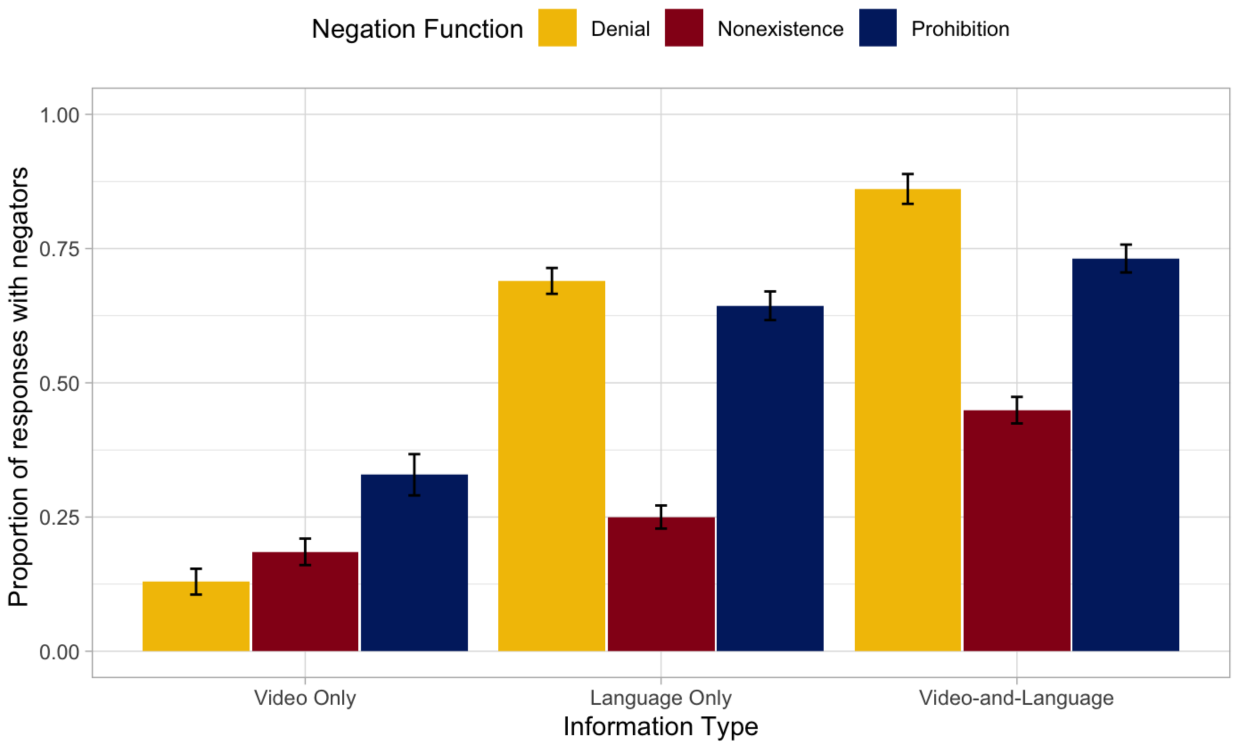

Figure 2: Proportion of negators included in participant responses for vignettes in which the target utterances included a negator by communicative function of that word (Denial, Nonexistence, Prohibition) and the amount of information provided (Information Type) along the x-axis. Error bars indicate ±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 (S1 in Appendix 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.

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.

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.

Methods

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.

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.

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.

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.

[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.

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 Appendix A) 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, 2018), 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*. 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.

Stimulus Creation

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.

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*.

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.

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.

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).

Effects of Information Type

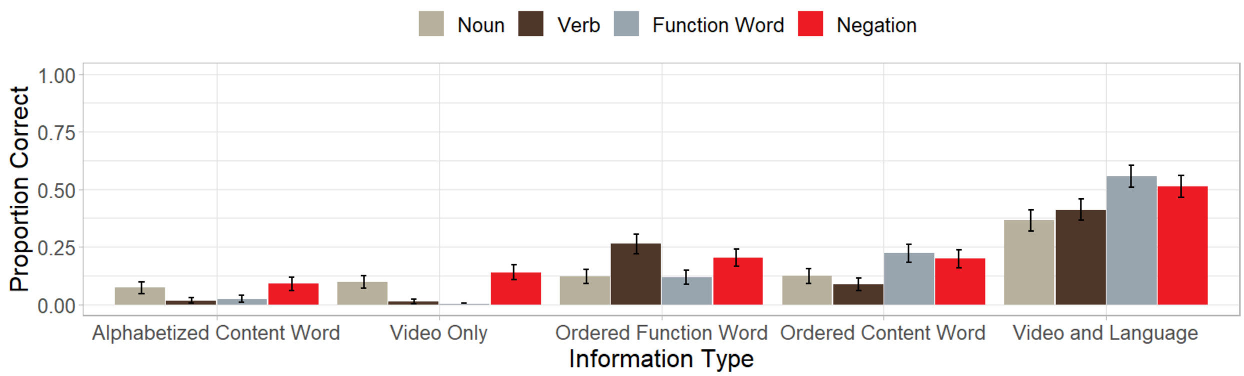

Figure 3: Proportion of correctly guessed target words, by Word Type and Information Type. All conditions included video. Error bars indicate ±1 Standard Error.

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 did not 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.

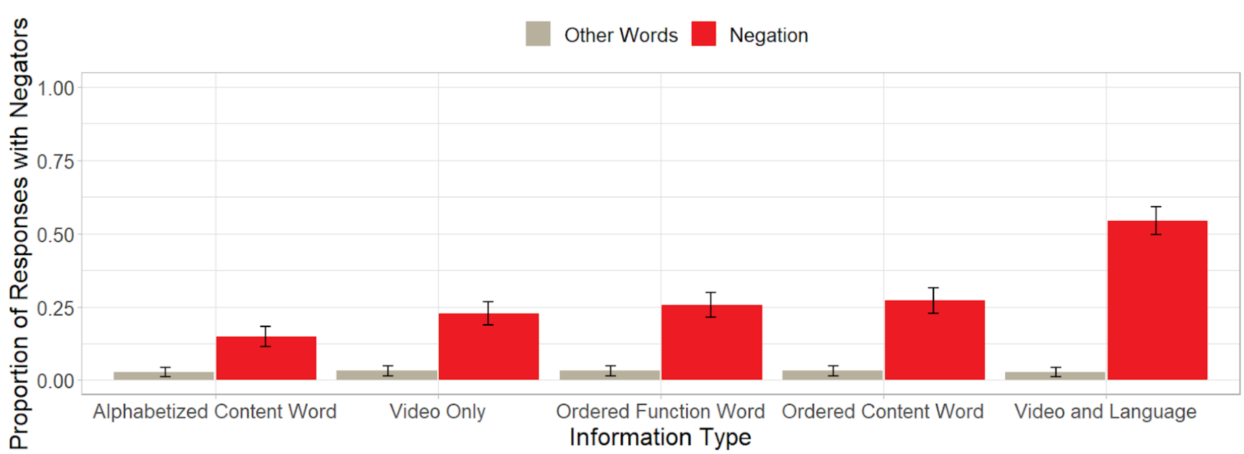

Figure 4: Proportion of guesses that were of any negation word (*no*, *not*, or any word ending in *n't*), as a function of whether the target word was actually a negation (*no* or *not*) versus any of the other word types (nouns, verbs, or function words), broken down by the amount of information. Error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

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.

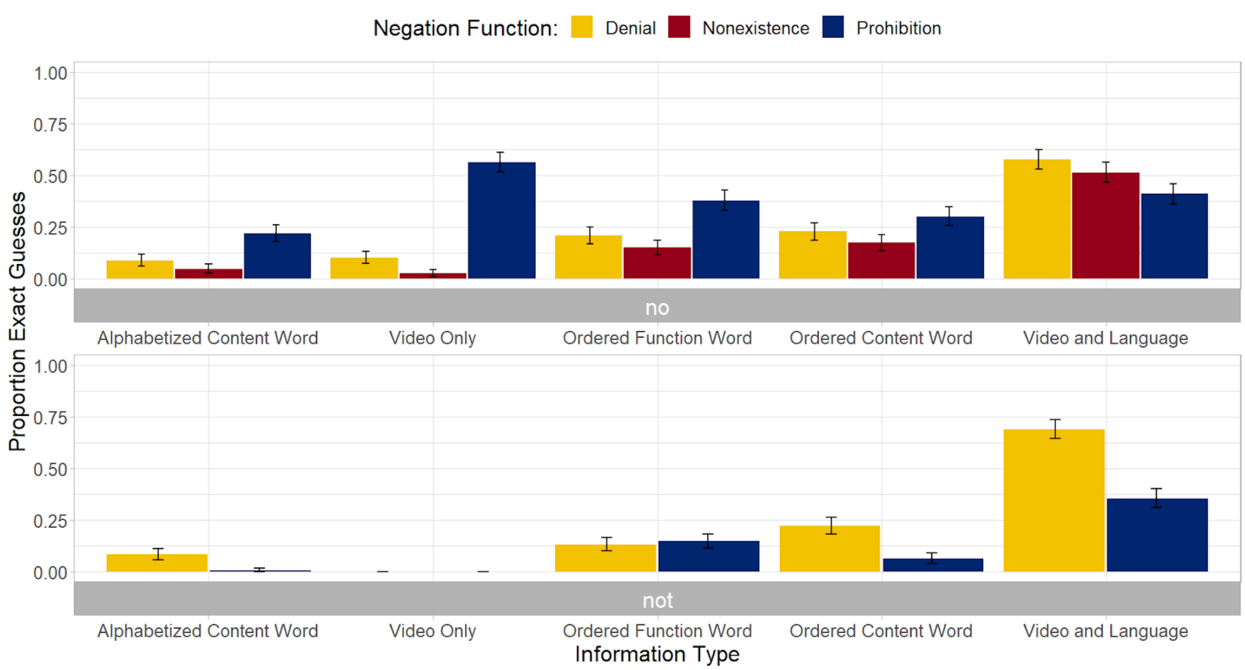

Figure 5: Proportion of correctly guessed target words, by Negation Word (above: *no*, below: *not*), the communicative function of that word (Negation Function: Denial, Nonexistence, or Prohibition), and the amount and type of language information provided

(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, ($\beta=0.69$, $z=4.79$, $p<0.001$, OR=2.0) and Denials were in turn easier to identify than Nonexistence uses ($\beta=0.40$, $z=2.51$, $p=0.01$, OR=1.49). In contrast, when the target word was *not*, Denial uses were the ones that were easier to identify than Prohibitions ($\beta=3.15$, $z=6.29$, $p<0.001$, 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 ($\chi^2(2)$=35.14, p<0.001) and Information Type ($\chi^2(1)$=37.79, p<0.001) were both significant, as was the interaction between them ($\chi^2(2)$=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).

[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.

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* 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 is not 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.

General Discussion

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.

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 et al., 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.

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.

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 Appendix B).

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.

Denial

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?

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, 2018; Reuter, Feiman, & Snedeker, 2018; Nordmeyer & Frank, 2015; 2023), 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.

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.

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.

CHAPTER 2: FAILURE IS ALWAYS AN OPTION
EVENT PERCEPTION & NULL-DISCOURSE NEGATION

"All of old. Nothing else ever. Ever tried. Ever failed.
No matter. Try again. Fail again. Fail better."
- S. Beckett (1989)

Truth-functional negation, which reverses the truth-value of a proposition (Horn, 1989), poses a remarkable challenge to accounts of word learning. Imagine you are a language learner, who has not yet acquired the terms for truth-functional negation in your language (e.g., *not*). When someone utters, "It is not raining," what information is available to you from observation of the referent world to help you learn what *not* might mean? Whenever it is not raining, it is by definition any number of other things (e.g., sunny, snowing, cloudy), and these other things are positively observable in the world. If young language learners are limited to the "here and now," how can they attend to the *not,* which serves as the invisible background to the here and now? Despite these theoretical difficulties, there is recent work suggesting that children are capable of understanding sentences with negation as early as 18 months of age (De Carvalho et al., in press) but certainly by 24-30 months (Feiman et al., 2017; Austin et al., 2017). Regardless of the exact age at which children acquire negation, it is remarkable that toddlers learn the meaning of such an abstract word when they are still building a basic understanding of the vocabulary and syntax of their language, have relatively limited Theory of Mind (Wellman, Cross, & Watson, 2001) and inhibitory control abilities (Lagattuta, Sayfan, & Monsour, 2011; Merkley, Thompson, & Scerif, 2016; Wright et al., 2003), knowledge and skills that could in principle be necessary for negation use and understanding (Lasnik, 1972; Cuccio, 2012; Nordmeyer et al., 2016).

Given that children master negation despite these limitations, it suggests they might actually be learning the meaning of, e.g., *not* via "observation," connecting utterances containing negation to their perception of the immediately surrounding referent world, even in a relatively egocentric way. But, how could this be? We offer an account and in-principle demonstration using adult and child participants. We propose that event perception can spontaneously provide a construal of an event (e.g., hitting) to an observer before the event completes, and so cases of failure (i.e., unfulfilled event expectation) may provide strong and stable evidence of a negative construal of an event. When failures are commented upon, negative construals are equally accessible to interlocutors jointly attending to such an event and as such they would be particularly helpful learning instances for negation. Here we ask can a language learner, when not given access to potentially helpful linguistic information to infer negation (e.g., the co-occuring words, the syntax of the utterance), nevertheless infer the presence of negation in that utterance? Or do they need this additional linguistic evidence to infer negation? If the former is true, it points to an important new line of inquiry regarding early acquisition of truth-functional negation. Thus we aim to investigate how a language learner (child or adult) might gain an understanding of truth-functional negation based on information that may be available in the circumstances of its use. Documenting when one expects negation to describe the here-and-now with adults would surely be a useful first step. By artificially generating such situations rather than starting with the naturally occurring ones, we are able to target those circumstances that we believe are theoretically relevant, thereby giving us some initial assurances about the feasibility of our account.

When considering the perception of events and their failures, it is important to note that certain early-arriving situational features may provide a given categorization of an event (e.g., either at an abstract level or description like VERB OF BOUNDED EVENT, or more fine-grained like HIT), before this event completes. Thus, as we watch a baseball player preparing to bat, perception furnishes observers with the expectation that hitting will occur, but as the player misses entirely, we are left with an unfulfilled expectation. Expectations about an event, when foiled, could be particularly promising learning instances for negation. This proposal builds on observations that failure (Choi, 1988) or unfulfilled expectation (Pea, 1980) are contexts in which children can at times produce *not.*This is because failure takes advantage of and is evoked by the immediate context (much like nonexistence; Choi, 1988). We explore this potential here with a series of experiments on children and adults using the Human Simulation Paradigm (Gillette et al., 1999), in which individuals attempt to guess aspects of a message simply by observing their circumstances of use. This paradigm allows us to examine whether in principle perceived circumstances of failure encourage the inference of negation in a message describing such events, and whether linguistic evidence provides further assistance. We will first briefly touch on the perception of events, and children's understanding of negation, and the learning of words by observation before turning to the details of the study.

Grounding negation in a perceptual process

Recently there has been a renewed interest in whether and how individuals perceive absences (e.g., Farennikova, 2013; Sorensen, 2008; Soteriu, 2011). This question is a natural extension of the longstanding debate surrounding the ontological status of negative facts. In the words of Horn (1989), "If an affirmative statement or judgment is about some positive fact in the world, what on earth can a negative statement or judgment be about? To what negative fact, if any, does it correspond or refer?" Similarly, what perceptual evidence grounds the experience of absence? An absence of an item will lead to the perceptual input of whatever is present (e.g., if the fur on a bear is missing, the light will reflect from its skin), so it's unclear how it could give quick insight into what is not. This might suggest that we determine absence by elimination, perception provides us with those objects that are present, and we search amongst them (and fail). However, some have suggested that this does not adequately reflect the phenomenology of absence, that we do in fact have a moment of recognition of absence which is immediate, and thus too quick to be the result of checking against all present items (Sartre & Barnes, 2003). As a result, the role of expectation in perceiving absences has been of critical interest (Farennikova, 2013; Sorensen, 2008; Soteriu, 2011). Such approaches may gain considerable attention due to increased general interest in predictive processing (Friston, 2010; Friston &

Kiebel, 2009; Keller & Mrsic-Flogel, 2018). By turning our focus away from objects (and thus absences) and instead to events (and thus failures), we aim to capture key insights from the philosophical literature (i.e., the importance of expectation, and violation) while exploiting the growing literature of event perception. A perceptual processing basis for unfulfilled expectation could lie in a top-down influence exerted by expectation (i.e., an object template during a visual search as in Farennikova, 2013), but events allow us to consider cases in which recognition and verification do not occur at the exact same moment.

There is growing evidence that perception is sensitive to several higher-order categories, including animacy (van Buren et al., 2015), the type of scene (Oliva & Torralba, 2001), as well as spatial, causal and event relations (Biederman, Mezzanotte, & Rabinowitz, 1982; Rolfs, Dambacher & Cavanagh; 2013; Hafri, Trueswell & Strickland, 2018; for review see Hafri & Firestone, in press). If something were to be recognized before it could be fully verified, then this could provide a clear case of an expectation (developed during recognition) going unfulfilled (during verification or completion). Events serve as particularly fertile ground for such an approach as various experiments have shown that adults (and children) are capable of categorizing or recognizing events prior to their completion. Work in this area by Hafri et al. (2018) has found that when performing a task unrelated to judging the roles of actors in an event (such as clicking on the person wearing the red shirt, or clicking on the girl), there was a reliable increase in reaction time when the roles of those actors changed. Since no part of the instructions encouraged attention to the actor roles, the authors argue that this indicates that event role information is extracted spontaneously from scenes of events. In an experiment aimed to establish an effect of recognition of event class rather than event roles, Ji & Papafragou (2020) showed that adults could successfully learn distinctions between videos of events which had a clear end state and those which did not, even in a condition where they were not shown the end state. These studies suggest that adults develop expectations about events as they unfold. This may be due to event perception (Hafri et al., 2018) or event cognition (Ji & Papafragou et al., 2020, though event perception could play a role as well, since unbounded events are repetitive (Strickland et al., 2015)).

Since events regularly fail in a way that objects do not, and there is reason to believe observers develop expectations about events before their completion, failures could serve as a clear microcosm for unfulfilled expectations. There may be some events or some class of events that have stable perceptual features which facilitate this predicted recognition of events. Predicted recognition is a better term for present purposes than *recognition* as it is only once an event has completed that it may be recognized by an observer as belonging to some class. Through this lens, a failure to catch is meaningfully **not a catching event**, thereby making a contradictory ("He did not catch the ball") accessible as a possible description, as opposed to a contrary which may lack overt negation ("He missed the ball."). If this is the case, then cases of failure may be particularly helpful for a language learner in mapping the semantics of negation onto a word, since it could make a negative construal easily accessible in a reliable way. Failures may therefore serve as reliable "null-discourse" contexts that require minimal set up or tracking of information. This is important, as negation has often been discussed as used when the state of the world is at odds with one's expectations (see Horn 1989, p.180-183 for a related discussion).

Such an account would be desirable as it provides: 1) a difference between violation of expectation (a surprise party) and unfilled expectation (no one showing up to your regular party); 2) a perceptual process that reliably supports the inference of negation that could presumably be exploited in language learning by providing individuals (with similar perceptual systems) a means by which a shared belief can manifest. And, importantly, it is therefore agnostic as to whether absence of failure is literally perceived, and thus is compatible with views of cognition that allow for top-down effects as well as those that do not.

Comprehension and production of negation in childhood

There is good reason to believe that young children learn to pair *not* in English with truth-functional negation without any explicit training, instead coming to an understanding of this mapping from everyday interaction with their parents. Comprehension experiments have found that children perform above chance in tasks that require understanding negative statements (e.g., “The ball is not in this bucket.”) once they are around two-and-a-half years old (Austin et al., 2014; Feiman et al., 2017). Using a habituation paradigm, de Carvalho et al. (in press) found that French speaking infants (18mo) were able to understand negation as they responded in a manner consistent with understanding the meaning of negation (*n’-pas*). In addition, child production data has found that children begin using negators to Deny the truth of a statement around this time (e.g., “That is not a cat.”;Bloom, 1968; Choi, 1988). Before this age, negators are typically used to Reject (“No veggies!”) and comment on Nonexistence (“No daddy?”). This taxonomy of negation was based on communicative function (Rejection, Nonexistence, Denial) and was developed to explore the earliest productions of negation (Bloom, 1968), and was was expanded by Choi (1987) and both this and other work (Pea, 1980) have suggested that utterances which describe Failure (or, more generally unfulfilled expectations) may be less cognitively taxing than Denial. This is because failures just require a single expectation to be tracked as opposed to keeping some aspect of the discourse, or some other interlocutor’s expectation or knowledge state in mind. Like Denial, however, such Failure uses tend to be expressed with typical sentential negation (e.g., “He did not jump”), so failures may serve as a more reliable learning instance for negators early on.

Taken together, the ability to comprehend and produce truth-functional negation in one’s language is present in children prior to two and a half years of age. This means most children acquire it before they ever experience any form of schooling. It’s more likely, then, that like other words and aspects of language learned at those times, that they are learning it implicitly from observing the world around them and the language used to refer to it. The observation of failures may be an important case in this regard.

Learning by observation and beyond

How might such a proposal work and what other sources of evidence might need to be available to a learner to support learning of negation in this way? To do this, we need to understand what is already known about how the observation of the referent world contributes to word learning and how it interacts with developing linguistic competencies. Observation of the referent world may assist in the acquisition of at least some words. There are many challenges to making such a proposal, and further nuance is required, but it seems clear that speech that refers to the here and now (i.e., the immediate perceived environment) could be quite useful (Locke, 1690). Of course, if the world were all a language learner could depend on, then it's unclear how exactly they could acquire the meaning of words in their language that do not correspond neatly to the physical world, like *think*. Fortunately, developing linguistic competencies may allow learners to exploit information present in an utterance to assist in further word learning (Gillette et al., 1999; Gleitman, 1990; Gleitman et al., 2005). As outlined by Gleitman and colleagues, there are in principle multiple evidentiary sources available to the individual when learning a word, and certain vocabulary items may depend on one source of evidence more than another with these differences in information necessity explaining which words are learned before others. For example, a language novice who is learning their very first words, likely only has access to the immediate situational context to infer word meaning, leading to the early acquisition of concrete terms that are most often used to refer to the here and now (the so-called "easy words", Gleitman & Trueswell, 2020). Once armed with this seed vocabulary, a learner may use the occurrence of these words to learn other words within a sentence. For example, a child who already knows the word for *cookie* will have a better chance of learning the meaning of *eat* when hearing someone say "She eats the cookie." And a slightly more sophisticated learner, who has access to syntactic information could determine the meanings of more abstract terms, such as the fact that give and tell are acts of transfer (e.g., "John told a story to Mary"), in a process known as syntactic bootstrapping (Gleitman, 1990).

Some initial evidence for this "information change" account of vocabulary growth came from a study of adults using the Human Simulation Paradigm (HSP) (Gillette et al., 1999), a method we take inspiration from in the present production experiments. The logic of the HSP is as follows. If children's early vocabulary is best understood as a function of their ability to exploit the evidentiary sources available, then adult speakers should be only as good at guessing the meanings of novel words as children when limited to the same sources of information. If, however, children's early vocabulary is best understood as a function of their conceptual ability, then naive adults should be better at guessing the meanings of novel words than children are when placed in a similar position. In this approach, adults "simulate" children at various stages in their linguistic development by being given access to different parts of the information present in natural parent-child interactions. These naive observers were tasked with guessing what the parent had said in the video, which provides an estimate of whether a learner could in principle guess that the message contained a given word. In one condition, designed to mimic early learning by visual observation only, participants watched a brief (40 sec) video of a natural parent-child interaction which is silent other than a single beep and are tasked with guessing what the parent uttered at the beep (either a noun or a verb). Naive observers were much more accurate at correctly guessing nouns than verbs, when they were provided with only the scene itself. Importantly, naive observers showed higher accuracy in identifying verbs when they were given additional linguistic information from the parent's actual utterance. For instance, in an attempt to 'simulate' a more advanced learner who knows the meanings of some nouns, different observers were supplied not only with the video but also an alphabetically ordered list of any nouns used by the parent for that utterance. This resulted in a modest but reliable boost in verb accuracy, suggesting that knowledge of co-occurring content words may facilitate verb learning. Moreover, observers who were given access to the syntactic information from the parent's actual utterance (where content words were replaced with nonsense words that are phonotactically permissible, as in "The blicket GLORPS that the dax is felmish.") were even more accurate (Gleitman et al., 2005). This means that structural information from a sentence helps a learner identify the meaning of abstract terms, such as many verbs (e.g., think, know). This account has been supported by numerous other HSP studies (Snedeker & Gleitman, 2004; Yurocvsky et al., 2013; Fitch et al., 2021), including negation (Gomes et al., 2023) and suggests the HSP provides a useful guide for how information accessibility shapes vocabulary growth. Though most HSP studies use adult participants, it has also been conducted with children (Piccin & Waxman, 2007).

While the HSP suggests access to newly developing sources of information can help, the paradigm does not, on its own, speak to the contribution of observation itself. Some additional work inspired by the HSP attempts to address this question, and is of particular interest to the present work. In addition to the effect of syntax, work by Papafragou et al. (2007) found that false belief contexts in particular may increase the “observability” of *think*. These contexts were ones where actors attempted to complete a routine action (use a pitcher to pour water) and had unintended effects occur (no water came out, other things came out). They presented children (3-5yo) and adults with false-belief scenes and asked them to guess the meaning of a word in a statement describing the scene with either credal syntax (sentential complement; e.g., “Matt GORPS that his grandmother is under the covers”) or with transitive syntax (e.g., “Matt GORPS a basket of food!”). They found that when scene and linguistic evidence converged (i.e., false belief scene with credal syntax), children and adults were significantly more likely to produce credal verbs than in any other condition (41% and 85% of responses respectively). This work with credal words like *think* or *belief* indicates that despite their abstractness there may be particular situational contexts in which their presence may be more easily inferrable when they are capable of exploiting linguistic evidence. Since negators in English occur in syntactic frames which are linearized in ways that are consistent with other kinds of words (e.g., quantifier for *no*, adverb for *not*), it is unlikely that linguistic evidence alone will encourage the inference of negation in a message. However, a serendipitous pairing of linguistic evidence with particular contexts (e.g., failure) may provide an ideal learning scenario for *not*.

Current Project

We present a series of studies designed to examine whether failures encourage the inference of negation in a description, and whether this interacts with linguistic evidence. We are interested in 1) whether failures may serve as a null discourse context for inferring the use of a negator in a description or 2) whether access to additional linguistic information is still necessary in such cases to infer the use of a negator. In all three experiments in Study 1, adult participants were presented with videos that depicted actors engaged in simple actions and asked them to guess how parents described these scenes to their children. We specifically designed pairs of videos that minimally differed along the lines of theoretical interest. A Success video depicted an actor completing an anticipated event, while a Failure video depicted an actor failing to complete an anticipated event. For example, in one video, the actor successfully opens a door, and in the matched Failure video the actor struggles (and fails) to open the door. Our prediction was that Failure videos would elicit significantly more guesses that include negators (*no*, *not*, and *-n't*) than the matched Success videos.

In addition, each experiment in Study 1 compared this base rate of hypothesizing negation against participants who observed the very same video but were also provided with additional linguistic clues that may support the hypothesis that a parental description of said video would include a negator. We created these linguistic "hints" from naturally occurring utterances from other subjects who had already described these videos in an earlier norming study. Each experiment was designed to offer observers different classes of linguistic clues that could theoretically support a guess that negation was present, at least when combined with the related scene.

With the results of Study 1 in hand, Study 2 turns to the question of how child observers might responds to these same events. The focus in Study 2 will be on the extent to which scene information alone can provide clues to negation for children. Study 2a asked 20 child participants (4-7 years) to describe the same set of failure and success videos as adults did in Studies 1a-c. Study 2 did not include conditions with linguistic hints, as the central question it seeks to investigate is whether failures can also serve as null discourse context for young children. In Study 2, both child and adult participants watched both Success and Failure versions of each event in two separate blocks. Participants were assigned to either watch all the Successes First, or the Failures First. As elicitation tasks for negators often involve repetition to set an expectation, the Successes First condition was intended to ensure participants had an opportunity to develop expectations about these events within the context of the task. Our prediction is not about whether ordering affects negator production, but rather that an effect of Event Type should remain regardless of Order of Exposure.

Taken together, we found that our predictions about events fail, we negate more often than when they do not. The adult data suggests that, in these cases, linguistic evidence is only helpful when the evidence is relatively rich and includes both syntactic and semantic information from the linguistic context; syntactic evidence alone, or rudimentary semantic evidence alone, did not increase negator production. The child data additionally suggests that this is seemingly not dependent on later improvements to e.g., memory, inhibitory control, world knowledge. This suggests that a learner capable of pre-recognizing events is capable of inferring the presence of negation in a message, and that this ability may improve as they develop a more sophisticated understanding of their language.

Study 1: Contributions of Linguistic Knowledge

In Study 1a, we presented observers with either the scene alone, or the scene plus a list of content words in alphabetical order (which, as in other conditions, were created from the typical descriptions made by participants in a norming study). The logic here is that providing an observer with the visual scene alone without any linguistic support simulates a learner at the earliest stages of acquiring linguistic meaning, i.e., a learner who does not know the meaning of any words; Supplementing the visual scene with the content words found in the utterance is intended to simulate a learner who has knowledge of some words in their language, but is still not able to take full advantage of the syntactic structure of an utterance. In Study 1b, participants were similarly either shown videos alone, or were additionally provided with Jabberwocky versions of the descriptions where content words were replaced with nonsense words, and negators were replaced with a blank. This experiment was intended to probe whether syntactic structure in the absence of lexical knowledge is able to assist in the inference of negation. In Study 1c participants watched videos alone, or with the full-linguistic hint with a single word missing. This final condition is intended to explore whether fairly sophisticated linguistic knowledge is necessary to meaningfully assist in the inference of negation. Or, put another way, whether a learner who has knowledge of the structure of an utterance and the meaning of all the words but one, would be able to infer the meaning of a negator with greater ease. Of course, as a consequence, developing linguistic knowledge increases the probability that any given utterance will fit such a description.

The pre-registration for all three experiments can be found on OSF at the following link: https://osf.io/wt6p2.

Study 1a: Content Words

Participants were shown a series of videos of common events and asked how parents described them to their children. In one condition, they received the video alone, to put them in the shoes of an early language learner that has to rely solely on visual observation. In the other condition, they were provided with a list of content words they were told were used by the parents in describing the scene, which was intended to correspond to an early stage in word learning, where some words may be known, but the grammar is not.

Method

Participants

Thirty-seven English-speaking adults were recruited from MTurk to complete the task, with 12 being assigned to each condition (i.e., either Video Only, Video with Language) and 12 being assigned to a control condition (Language Only). Participants were excluded if they provided responses that indicate they did not understand the instructions (i.e., not providing a sentence) or that were ungrammatical more than 20% of the time. This resulted in one exclusion, leaving 36 participants for data analysis.

Materials

Visual Materials

Forty-four short videos of actions under 20s each were developed for the following experiments. Twelve actions were used to construct the experimental items, with each action being filmed as a success (e.g., lighting a match) or a failure (e.g., striking a match against a matchbox without it lighting). This resulted in 24 experimental items, with 12 successes and 12 failures. Twenty videos were developed for the filler trials, so that each participant saw 32 videos total were presented to each subject (12 experimental videos (6 failures and 6 successes) plus 20 filler videos). The experimental videos are available on the paper's OSF page, at https://osf.io/y7j3w/.

Linguistic Materials

A norming study provided the grounds for the linguistic materials used throughout the experiments. A different group of adult native speakers of English (N=68) than the ones tested in the experiments reported here provided descriptions of the videos on their own. We noted the objects, predicate, tense, and, in the case of negatives, the form (e.g., *no* vs. *not* vs. *-n't*) and the auxiliary used most frequently to guide the design of an "idealized" description for each video. We noted the most common features of negative descriptions of a given video to construct an "idealized" version that shared the most common features. Participants in the Linguistic Information conditions of Study 1a were provided with the content words from these descriptions listed in alphabetical order (so as to not provide potential information to syntax). For example, the idealized description for the failure version of the match-lighting vignette was “A person can’t light a match,” and so the Linguistic Information presented to participants in Study 1a for this vignette would have been “light, match, person.” The corresponding Linguistic Hint used in each trial for each Experiment is listed in full in Appendix C.

Procedure and Experimental Design

Participants took part in the experiment over the internet using the PCIbex experimental platform (Zehr & Schwarz, 2018). All participants provided informed consent and written instructions pertaining to the task before beginning the experiment.Participants were randomly assigned either to the primary experiment or to the “Language Only” condition.

For those assigned to the primary experiment, the procedure was as follows. Participants viewed 32 videos that they were told parents had previously watched and had described to their children. Participants were tasked with guessing how the parents described them. At the start of each trial, the video would begin playing automatically. Each video played only once, and participants were prompted for their response after the video ended. Once the video ended, a text box would appear for them to type their answer into. Those assigned to the Video with Language condition were also provided with the ‘linguistic hint’ that would appear below the video and above the text box. The hint consisted of an alphabetized list of the content words from the idealized sentence (see above). Participants in this condition were told to use all of these words in their response but could also add any additional words if they wished. The instructions at the start of this condition were as follows:

> You are about to see a series of silent videos and a list of words from sentences that English-speaking parents have used to describe these videos to their children. After each video, you will be presented with a list of words (in alphabetical order) that parents used in describing the video to their children. For example, after watching a video you may see a list like "dad, doggie, feeding"

> and your task would be to provide a sentence using these words. In this example, perhaps "Dad is feeding the doggie," would work, or another sentence using those words. For each video, we want you to provide a sentence that you think a parent might have used. Once you submit your answer, the next video will begin. <u>Videos will only play once</u>. A list of words and text box will appear once the video finishes playing. In the box, type a sentence you think a parent might have used to describe the videos. Please do your best to guess what the parent might have said.

Those in the Video Only condition received the same instructions, excluding any mention of alphabetically-ordered words. They were then provided no hint and freely produced their own sentence. Once the participant typed in their answer and clicked submit, the experiment proceeded to the next trial. Each participant saw all twelve experimental videos, but saw only the success or failure version of each one (6 failures, 6 successes). The order of trials, including fillers, were randomized for each participant.

Thus, Stimulus Type (Video Only vs. Video with Content Words) was a between-subject factor whereas Event Type (Success vs. Failure) was a within-subject factor, resulting in a 2x2 experimental design. Four experimental lists were used to implement this design. Two of these lists were for the Video Only condition, and two were for the Video with Content Words condition. For the Video Only condition, the first of these two lists was created by randomly assigning half of the trials to the Success video condition and half to the Failure Video condition. List two reversed these condition assignments. The same two lists were then used to create the Video with Content Words condition, except here the content words were provided. Each participant was randomly assigned to one of these four lists.

We ran a separate control condition to examine whether the linguistic information alone (i.e., the list of content words) presented any sort of bias in response. In this Language Only condition, participants were presented with the 32 content word lists used in the primary experiment. Participants were told the list corresponded to content words used by parents to describe a series of videos, but they were not presented with these videos. They saw one word list at a time, and were asked to guess the parent's full utterance (a sentence) using the words present in each list. They were free to add words if they wished. Trial order was randomized for each participant.

Results

Each participant response was binarily coded according to the presence of a negator (here defined as *no*, *not*, or *-n't* contraction). Figure 6 presents the proportion of responses that included a negator by Stimulus Type (Video Only, Video with Language) and Event Type (Failure vs. Success) for the conditions which included Video.[5]

[5]In the Video-and-Content-Words condition, the Success events had no negator responses at all (0%), which caused the generalized linear model (using the lme4 package in R) to fail. We randomly selected a single trial from a single subject and changed the binary negator response from 0 to 1.

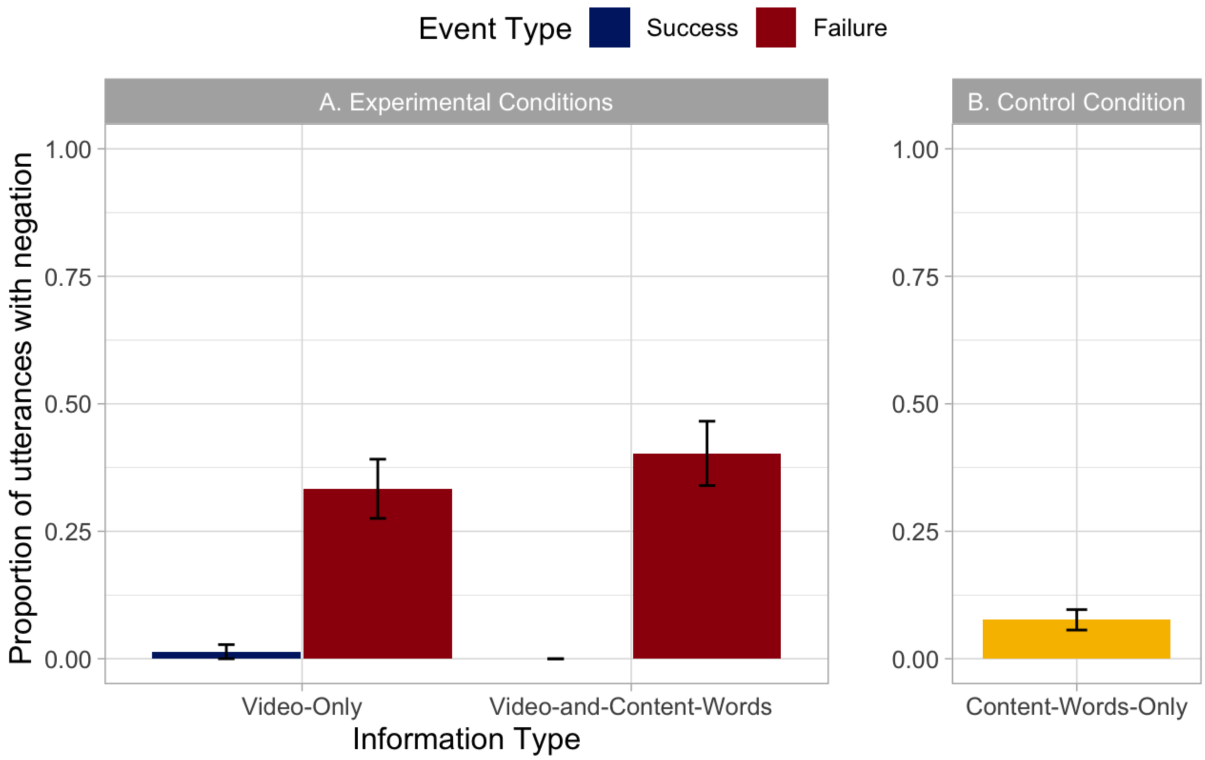

Figure 6: Proportion of participant responses which included negators for Video and Video with Language conditions. Study 1a. Data are presented by Event and Stimulus Type, with error bars indicating ±1 Standard Error.

As can be seen in Fig. 6, participants were significantly more likely to provide a negator in their response to failure scenes (36.8% of the time) as compared to success scenes (1.4%) when collapsed across Stimulus Type, and this effect was not significantly enhanced by the addition of Linguistic information (i.e., alphabetized list of Content Words). Those in the Language Only "control" condition, who did not see the video and only saw the list of Content Words, produced a negator only 7.6% of the time. These observations regarding the differences in means were supported by a multilevel logistic regression using glmer (Bates et al., 2015) to predict negator presence using Event and Stimulus Type and their interaction as fixed effects, and by-participant and by-item random slopes as random effects. We found a main effect of Event Type as we predicted, (β=4.014, SE=0.786, z=5.107, p<.001) but did not find an effect of Stimulus Type (β=-0.193, SE=0.83, z=-0.232, p=0.817), nor an interaction (β=0.774, SE=1.5514, z=0.499, p=0.618). Pairwise comparisons were run to compare performance to the control condition (Language Only), predicting negation with Event Type and Stimulus Type (in this case, Language Only either Video Only or Video with Language). These analyses found that only the two Failure conditions significantly differed from the base rate (Video Only: $\beta=1.705$, $SE=0.494$, $z=3.449$, $p<0.001$; Video with Language: $\beta=2.082$, $SE=0.679$, $z=3.063$, $p=0.002$).

Discussion

Our predictions about the effect of failure scenes encouraging negator responses was held, but there was no additional assistance of content words as we initially predicted in our pre-registration. Since content words allow for an event or object to be picked out, we predicted that this would encourage participants to describe the failure as opposed to other aspects of the scene, but this did not seem to be the case. Instead, participants in both the Video-Only and Video+Language condition produced negated sentences at similar rates. When they did not provide negation, they often described the intent rather than the event (e.g., “The person tries to light the match.”) or other aspects of the event (e.g. “The person is having trouble cutting the leaf.”). Describing intent was possible even in the Video+Language condition because participants simply added words like "tries to" (rather than *not*) to their response while still using all the hint words. A comparison of responses which included “try” between for Studies 1a-c is provided in Appendix D. Thus, despite providing an additional hint as to what is being discussed in a scene, content words did not adequately constrain the constructions in which those words might appear. In hindsight, this finding (of no effect of Stimulus type) is perhaps not surprising as it is well-documented that negation interfaces regularly with syntax, and as a result, syntactic information may adequately constrain participant responses to exclude these alternative responses.

Study 1b: Syntax

Study 1b was identical to Study 1a except for the linguistic hint that was used in the Video with Language condition. Here the "hint" was in the form of "Jabberwocky." In particular, each typical description of each video (generated from the norming study) was modified, such that all content words were replaced with nonsense words, so that participants would be able to make use of the syntactic structure without knowing the meaning of the words involved. So, if the idealized sentence had been "A woman is putting on lipstick," participants would have seen, "A tive is ___ on flous." This is intended to simulate a hypothetical learner that may have a fairly robust understanding of the syntax of their language, but is encountering a phrase with no known content words.

Method

Participants

A new set of thirty-eight English-speaking adults were recruited from MTurk to complete the task, with 12 being assigned to each condition (i.e., either Video Only, Video with Language) and 12 being assigned to a control condition (Language Only). Participants were excluded if they provided responses that indicate they did not understand the instructions (i.e., blank answers, providing multiple words) more than 20% of the time. This resulted in two exclusions, leaving 36 participants for data analysis.

Materials

Visual Materials

The visual materials used in the three experiments were exactly the same (see Study 1a: Materials).

Linguistic Materials

As described in Study 1a, a prior norming study provided the ground for the linguistic hints. Participants in the Linguistic Information conditions of Study 1b were provided with a

Jabberwocky version of the idealized description by replacing the content words from the description developed from the norming with nonsense words. In the case of contractions, auxiliary and negator were separated, and the negator was replaced with a blank. For example, the idealized description for the failure version of the match-lighting vignette was “A person can’t light a match,” and so the Linguistic Information presented to participants in Study 1b for this vignette would have been “A tive can ___ glizz a dax.” The nonsense Jabberwocky words were partly derived from previous experiments (Gillette et al., 1999; Snedeker & Gleitman, 2004) as well as the ARC Non-word Database (Rastle, Harrington, & Coltheart, 2002). For filler items, other parts of speech were randomly selected to be removed (either adjectives, nouns, adverbs, verbs, prepositions, and conjunctions).

Procedure and Experimental Design

The procedure and experimental design were the same as Study 1a, except for the following changes. For the conditions which include linguistic hints (Video plus Language, and the Language Only conditions), rather than a alphabetic list of content words, the linguistic hint was instead a Jabberwocky sentence with a fill in the blank (e.g., “A tive can ___ glizz a dax.”). Participants were told this was the parent’s utterance with all content words replaced with nonsense words. They were instructed to fill in the blank using English to complete the sentence. To make sure participants in these conditions understood their task, the instructions at the start of this experiment were as follows:

You are about to see a series of silent videos and sentences that English-speaking parents have used to describe these videos to their children. These sentences have been translated into a novel language. We have removed ONE WORD from the sentence. For each video, we want you to provide the missing word that you think a parent might have used. For example, after seeing the video, you might see a sentence like "The ___ dax is glorping." Your task would be to fill in the blank with a word in English that you think could go there. In this example, perhaps "yellow" or "happy" could work as in "The <u>yellow</u> dax is florbing.” Remember to provide your answer in English and not in the nonsense language! Once you submit your answer, the next video will begin. Videos will only play once. A sentence and text box will appear once the video finishes playing. In the box, type ONE WORD. Please do your best to guess what the parent might have said.

The Video Only condition was identical to Study 1a.

Results

Responses were coded for negators as described in Study 1a: Results. Participants in Study 1b performed similarly to those in Study 1a (see Fig. 7). Participants produced more negators in their responses when describing a Failure Event (38.9%) than when describing a Success Event (5.6%) when collapsed across Stimulus Type, and there was similarly no effect of Stimulus Type on the production of negators. In the Language Only control, participants included negation in only a small number of responses (8.33%).

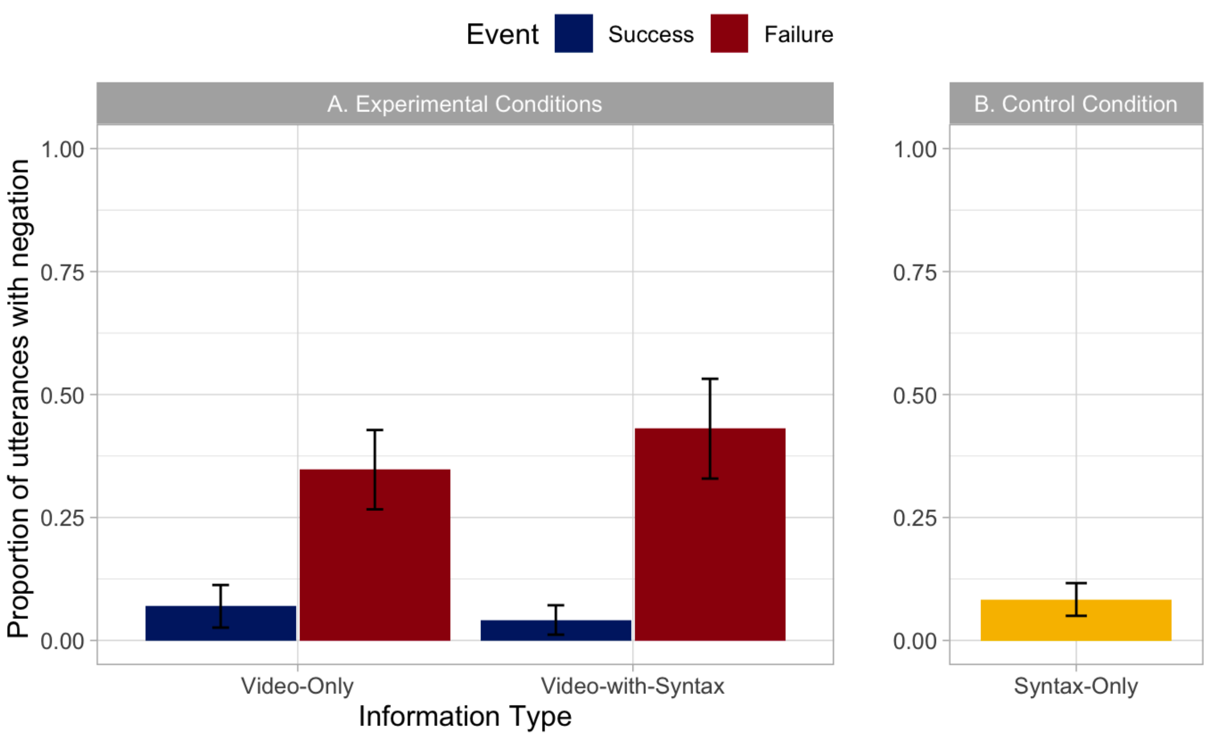

Figure 7: Proportion of participant responses which included negators for Video and Video with Language conditions. Study 1b. Data are presented by Event and Stimulus Type, with error bars indicating ±1 Standard Error.

We analyzed these results using the same model as described in Study 1a: Results. Just like in Study 1a, we again found a main effect of Event Type as we predicted, ($\beta=4.011$, $SE=1.501$, $z=2.672$, $p=0.008$) and did not find an effect of Stimulus Type ($\beta=-0.149$, $SE=0.915$, $z=0.162$, $p=0.871$), nor an interaction ($\beta=1.379$, $SE=1.487$, $z=0.928$, $p=0.353$). Pairwise comparisons of each Event Type to the baseline condition (Language Only) found that only the two Failure conditions significantly differed from the base rate (Video Only: $\beta=1.705$, $SE=0.494$, $z=3.449$, $p<0.001$; Video with Language: $\beta=2.082$, $SE=0.679$, $z=3.063$, $p=0.002$). Pairwise comparisons differed slightly from those in Study 1a, with only the Failure condition of the Video Only condition differing from the Language Only base rate ($\beta=3.679$, $SE=1.471$, $z=2.501$, $p=0.012$), though there was a trend when comparing the Failure condition of the Video with Language condition to the base rate ($\beta=10.93$, $SE=5.825$, $z=1.876$, $p=0.061$).

Discussion

Our predictions about the effect of failure scenes encouraging negator responses were replicated. Since syntax does not uniquely support a negator, we predicted in our preregistration that it would have no effect, which was confirmed. It could be because participants guessed that one of the nonsense words included negation implicitly (like *miss*) and hence did not need to add it in explicitly (Horn, 1989). Though there's no way to determine this, we do look at how often participants provided replies with implicit negation in Appendix D. Alternatively, they could have used an adverb prefixed with *un-* or *dis-*, but this only occurred once. Like content words, despite providing an additional hint as to what is being discussed in a scene, syntax did not adequately increase the guessing of negation.

Study 1c: Full Linguistic Information

Study 1c was identical to Studies 1a and 1b except for the linguistic hint that was used in the Video with Language condition. Here the "hint" was a full English sentence missing one word which was replaced by a blank. In experimental trials, the blanks were placed such that a negator could be grammatically inserted. If the idealized sentence had been “A woman is putting on lipstick,” participants would have seen, “A woman is ___ on lipstick.” This is intended to simulate a hypothetical learner that may have a robust understanding of their language, but is encountering an utterance with one unknown word.

Methods

Participants

Thirty-seven English-speaking adults were recruited from MTurk to complete the task, with 12 being assigned to each condition (i.e., either Video Only, Video with Language) and 12 being assigned to a control condition (Language Only). Participants were excluded if they provided responses that indicate they did not understand the instructions (i.e., blank answers, providing multiple words) more than 20% of the time. This resulted in one exclusion, leaving 36 participants for data analysis.

Materials

Visual Materials

The visual materials used in the three experiments were exactly the same (see Study 1a: Materials).

Linguistic Materials

As described in Study 1a, a prior norming study provided the ground for the linguistic hints. Participants in the Language conditions of Study 1c were provided with a fill-in-the-blank version of the description by removing one word from the description developed from the norming. The conventions set in Study 1b for fillers and contractions were followed. For example, the idealized description for the failure version of the match-lighting vignette was “A person can’t light a match,” and so the Linguistic Information presented to participants in Study

1c for this vignette would have been “A person can ___ light a match.” The same hint was provided for both the Success and Failure versions of this video.

Procedure

The procedure follows that of Study 1b, except that in the conditions which included Linguistic hints, participants were given full sentences with a blank that they were told corresponded to parental descriptions of a series of videos. The instructions at the start of this experiment were as follows:

> You are about to see a series of silent videos and sentences that English-speaking parents have used to describe these videos to their children. We have removed ONE WORD from the sentence. For each video, we want you to provide the missing word that you think a parent might have used. For example, after seeing the video, you might see a sentence like “The ___ doggie is sleeping." Your task would be to fill in the blank with a word in English that you think could go there. In this example, perhaps "yellow" or "happy" could work as in "The <u>yellow</u> doggie is sleeping.” Once you submit your answer, the next video will begin. <u>Videos will only play once.</u> A sentence and text box will appear once the video finishes playing. In the box, type ONE WORD. Please do your best to guess what the parent might have said.

Results

As in Studies 1 and 2, participants were much more likely to produce responses with negators for Failure Events than Success (see Fig. 8). But, unlike the previous experiments, participants were even more likely to produce negators when provided with additional linguistic cues.

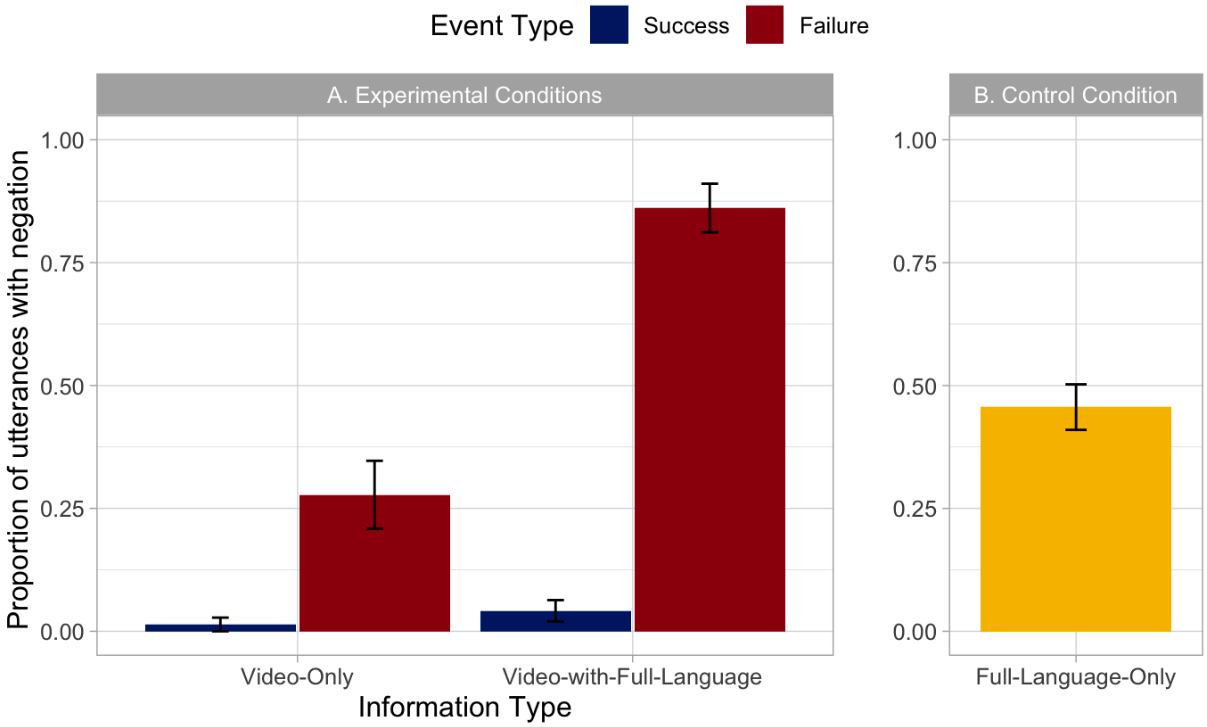

Figure 8: Proportion of participant responses which included negators for Video and Video with Language conditions. Study 1c. Data are presented by Event and Stimulus Type, with error bars indicating ±1 Standard Error.

Running the same analyses as were run for the other studies confirmed that there was a main effect of Event Type (β=5.52, SE=0.985, z=5.605, p<.001) and Stimulus Type (β=2.228, SE=0.918, z=2.426, p=0.015), with an interaction (β=3.735, SE=1.825, z=2.046, p=0.041). The base rate of producing negation in the Language Only control condition was much higher than in the previous experiments (45.6%). The pairwise comparisons were fairly similar to those in Study 1a, with a significant difference between the two Failure conditions from the Language Only control (Video: β=-1.208, SE=0.002, z=-501.364, p<.001) (Video with Language: β= 2.908, SE=0.968, z=3.004, p=0.003). There was no reliable difference from base rate for the two Success conditions. Though, unlike in the previous studies, the model only indicated a trend in the Video Only condition (β=-6.416, SE=3.675, z=- 1.746, p=0.081) and failed to converge in the Video with Language condition.

Discussion

Our predictions about the effect of failure scenes encouraging negator responses were replicated again, and our predictions about the effect of Full linguistic information and interaction with Event Type were supported. The reliable interaction found in this experiment between Event Type and Stimulus Type indicates that the linguistic information in our experimental items did not deterministically increase responses with negators and instead the rate of negation responses was modulated by the scene which it describes. This suggests that such a learner, who has access to full linguistic information, displays higher sensitivity than a learner limited to visual information alone, presumably resulting in increased confidence, though that is beyond the analyses provided here.

Pairwise Comparisons of Studies

Cross-experiment comparisons were performed using a glmer model which predicted negation using Experiment, Stimulus Type, and their interaction as fixed effects and with by-participant and by-item random slopes. This model was run on data from the Failure Events only and was conducted in a pairwise manner. As detailed below, these models revealed no reliable effects or interactions when comparing the partial Linguistic Information conditions to each other (Content words (Study 1a) and Jabberwocky (Study 1b)). However, there were significant results when comparing partial Linguistic Information conditions to Full Linguistic Information conditions as in Study 1c (Content Words to Full, Jabberwocky to Full).

The model comparing Studies 1a (Content Words) and 1b (Jabberwocky) found no main effect of Experiment (β=-0.04108, SE=0.38263, z=-0.107, p=0.914), no main effect of Stimulus Type (β=0.61744, SE=1.22175, z=0.505, p=0.613), and no reliable interaction between the two (β=-0.14199, SE=0.76551, z=- 0.185, p=0.853).

The model comparing Studies 1a (Content Words) and 1c (Full Linguistic Context) found a reliable main effect of Experiment (β=1.14023, SE=0.35250, z=3.235, p<0.001) and a reliable main effect of Stimulus Type (β=1.77007, SE=0.36678, z=4.826, p<0.001). However, the effect of Stimulus Type was found to interact reliably with Experiment (β=2.84034, SE=0.71805, z=3.956, p<0.001). This interaction reflects the fact that the (null) effect of Stimulus Type found in Study 1a was much smaller than the (reliable) effect of Stimulus Type found in Study 1c.

Finally, the model comparing Studies 1b (Syntax) and 1c (Full Linguistic Context) found a reliable main effect of Experiment (β=1.339, SE=0.543, z=2.469, p=0.014), and a reliable main effect of Stimulus Type (β=2.23286, SE=0.561, z=3.978, p<0.001). However, the effect of Stimulus Type was found to interact reliably with Experiment (β=3.298, SE=1.103, z=2.989, p=0.003). This interaction reflects the fact that the (null) effect of Stimulus Type found in Study 1b was much smaller than the (reliable) effect of Stimulus Type found in Study 1c.

Study 1 Discussion

In our task, we used adults to "simulate" earlier learners by asking them to guess how parents would describe a video. Unlike the HSP, these videos were not taken from everyday parent-child use but instead artificially generated to test the generalizability of failures (Gilette et al., 1999). But like the HSP, adults were provided with information on different aspects of an utterance thereby "simulating" a learner at different stages of information access. Adults in the Video-Only conditions of Studies 1a-c simulated an early language learner that can only infer what a parent is saying based on the referent world. Strikingly, we found that adult participants were much more likely to use negation to describe actors failing at some action than succeeding, even with no additional linguistic information, across three sets of experiments. This suggests that event failures could serve as null discourse contexts for negator learning, and motivates additional work in the nature of the expectations violated by these failures as well as their relationship to lexicalization and production. In addition to event outcome information from the scene, we included conditions where we provided participants with additional linguistic information about the sentences uttered by parents (but in fact were based on prior descriptions by other adults). Adults in the Video-and-Language conditions of Studies 1a-c simulated learners with access to different aspects of the linguistic utterance. We showed that co-occurring linguistic information can increase negator productions for failures, but only when this information is relatively complete, as in Study 1c (where all that was absent was a negator).

When partial linguistic information was given in addition to the Video, as in Study 1a (Content Words) and Study 1b (Syntax), there was no significant increase in negative responses when compared to viewing the same failure Video alone. For Content Words, this may be due to "try" responses or implicit negation, while for Jabberwocky it may have been due to the difficulty of the task, or guessing non-negative items like adverbs or quantifiers. In the first experiment, participants in the Language + Video condition were presented with a list of content words that were in the "parent's" utterance in alphabetical order. Since this would provide participants with a clearer idea of what was being discussed about the scene, as well as the framing (e.g., the list always included the name of the action and not "miss" or "fail"), we predicted that this might assist participants in inferring negation for failure vignettes. However, we saw no significant difference between this condition and the Video Alone condition. Instead of including negators, participants often described the actors as "trying to X," which suggests that the content words were not effective because they did not block different constructions that related to different aspects of the events (e.g., intent, desire).

In the second experiment, the linguistic information was instead a Jabberwocky version of the target utterance, with content words replaced by nonsense words that obey English phonotactics. These Jabberwocky sentences were constructed such that there was additionally a blank where a negator could be inserted to negate the predicate (e.g., "The blick is ___ daxing") for both the Failure and Success events (in filler conditions, the blanks replaced other syntactic categories). If the lack of effects in the content words version was solely the result of unblocked "trying to" constructions, then the Jabberwocky sentences should produce an effect. However, as we noted, negation does not occupy unique syntactic frames that would exclude other alternatives when linearized (e.g., quantifiers for *no*), and this is true of phrasal negation as well. In all our Jabberwocky sentences, for example participants could have guessed an adverb (e.g., "The blick is quickly/almost daxing") instead of *not.* As we predicted, the syntactic information provided by Jabberwocky sentences did not significantly increase the inference of negation in the Video + Language condition. It was not until the third experiment, which included the full utterance and a blank where a negator could go, that we found a significant difference between the Video Alone and Video + Language conditions. Participants in the Language Only control condition of this experiment produced negation at significantly higher rates than in the other experiments (48.95% as compared to 7.64% in Study 1a (Content Words), and in 8.33% Study 1b (Jabberwocky)). But, notably, when linguistic information was presented alongside observation, there was an interaction such that negation production for Success Events was extremely low, and that of failures was high. This indicates that though the linguistic information alone may introduce some bias, it does not solely determine the inference of negation.

We began by proposing a mechanism by which failure could serve as a reliable learning instance for negation as proposed by acquisitionists who have studied the form and function of negation in the past (Choi, 1988; Pea, 1980). Our proposal was motivated by recent work in event perception (Hafri, et al., 2018; Ji & Papafragou, 2020) and the philosophy of perception (Farennikova, 2013; Sorensen, 2008; Soteriu, 2011). Prior work has shown that children and adults are sensitive to higher order features of events prior to their completion in a way that is quick and seemingly immediate, so we suggested that this almost automatic prediction furnished by perception may regularly result in a mismatch between what is predicted (e.g., a telic event, like hitting) and how the event actually unfolds. This would offer a reliable perceptual mechanism that would make unfulfilled expectations more jointly accessible to interlocutors commenting on a scene, and thus make an otherwise individual and private experience, into a shared one. Importantly, however, this mechanism must also be available to language learners. Study 2 seeks to explore whether this is possible, and, therefore, whether event failures truly are gem-like contexts for negation (Gleitman & Trueswell, 2020).

Study 2: Negation in null discourse contexts

In Study 2a, children and adults participated in a child-friendly version of the Video-Only condition of Studies 1a-c. Their task was simply to watch and describe a series of videos. As mentioned in the introduction, children rarely produce sentential negation until they are past the two-word stage, and eliciting it can prove challenging (Klimacka & Brunger, 1999). A common task to elicit negation from children involves presenting children with a set of figures (e.g., horses), all of whom are engaged in an action (e.g., dancing) but one. The experimenter asks the child what each figure is doing to prompt a particular verb (e.g., "the horse is dancing"), with the sole non-participating figure saved for the last question (Thornton et al, 2016; Thornton et al., 2018; see Thornton & Rombough, 2015 for a non-action version). These tasks are commonly used in Developmental Language Delay (DLD) research and by those investigating the syntax/morphology of negation, and so they are not set up to answer the question at hand. For example, in these tasks, the target of negation is typically present in some manner (it is true of something else currently visible, e.g., other horses are dancing), presumably to reduce the memory load, while in the case of failures, the failed event itself is fleeting. More importantly, these tasks ensure the child has a particular category in mind (e.g., dancing) before asking the critical question, while our aim is how children can home in on the meaning of negation if one does not already know that the category in question is relevant. A simple consequence of this is that these tasks require more set up than simply asking people to describe a video, they require an experiment to provide scaffolding through a series of questions, whereas the present task the experimenter only asks one question - "what happened?" Because this question is open-ended, task demands are reduced, and we may gain more insight into how relevant negating is to children. And, because the present task requires less set up, it probes a situation which may be quite common: a parent commenting on a failure in front of their child.

However, we remained curious about the extent to which the elicitation of negation in our task requires availability of the successful alternative. Despite the fact that stimuli were created out of simple everyday actions, adults nonetheless have much more experiences with such actions than children. Therefore, children could be worse at recognizing such actions are being attempted prior to their success. As such, making the successful action available within the context of the task may be necessary for drawing attention to the failure. To test this potential, we blocked Failure videos and Success videos, and randomly assigned participants to different Orders of Exposure, either watching all the videos with Successes First or all the videos with Failures First. If children are less likely to see Failures as such without being recently primed with Successes (and are thus less likely to negate), participants in the Successes First condition should describe failures with negators more often than those in the Failures First condition. Additionally, there should be an interaction such that if there is any effect of Event Type it is not observed in the Failures First condition. However, if failure is indeed a true null-discourse for these learners, then an effect of Event Type should be observed for the Failures First condition as well. Additionally, this blocking of condition, rather intermixing trials by condition, has the benefit of better managing challenges related to children's limited executive function abilities that may make it harder for them to switch between negating and not.

A secondary analysis was planned and pre-registered to analyze alternative constructions. Motivated by findings in Studies 1a-c, we anticipated that children might provide alternative constructions to describe failures, like "try" (e.g., "she is trying to touch her toes."), as well as words which contrast with possible alternative descriptions. On the one hand, children may recognize the failures as such but not tend to use negators to describe them, and "try" may remain available and thus potentially even more common. On the other hand, children may be less likely to use *try* either because unlike adults they do not find the failures remarkable (and therefore comment on something else) or because *try* is linguistically complex in a way that may impact producing sentences which contain it, as well as learning it via syntactic bootstrapping. It is a subject-control verb and children can struggle with such verbs in certain contexts until 4-5 years (Chomsky, 1969; Kessel, 1970; Hyams & Orfitelli, 2017). Using "try" in a sentence therefore might require more linguistic competence, which the younger participants in our study are likely still developing.

Study 2a: Child Study

Children (4-7 years) were shown a series of videos of common events and were asked to describe them out loud. Unlike Studies 1a-c, participants watched and described both the success and the failure video for each event. In one condition, participants watched a block of the success videos first to make the intended action locally available. In the other, they watched a block of the failure videos first to determine whether they elicit negators in the absence of additional context. The pre-registration for Study 2a and 2b can be found on OSF at the following link: https://osf.io/bcgq9.

Methods

Participants

Twenty participants from 4 to 7 years of age (m_{age}=5.85) participated and contributed data to the results reported below. All participants were recruited from Philadelphia-area preschools and schools. Parents of participants provided informed consent prior, and participants were given a choice of a book at the end of the study (Cat in the Hat, Fox in Socks, The Little Prince). Two additional participants participated but were excluded for providing inaudible responses (both to the experimenter running the experiment as well as a coder who checked the recorded audio). The participants were assigned to the Failures first or Success first conditions.

Materials

The visual materials were the same as those in Study 1a, except that one event, “light,” was not used, as it depicted the author attempting to light a match and we did not want to encourage this type of reckless behavior in children. This resulted in 11 different events in Study

2 rather than 12. To accommodate the larger number of critical videos than the previous study, filler items were not included in Study 2.

Procedure

Following other similar open-ended production tasks with young children (Papafragou et al., 2007; Piccin & Waxman, 2007), participants were told that they would be playing a game where they would describe a set of videos. The experimenter told the child they had recorded these videos with friends, and needed the child's help describing them out loud after they had finished playing. Participants wore a wired logitech headset with a microphone while completing the task so that their responses could be recorded for transcription. They were told that as a reward for describing each video, they would see a cute image of a frog, and that at the end the experimenter would check if they got them all. They were additionally told that the screen would turn blue (i.e., the background's color would change from white to blue, with various yellow question marks) when videos ended, to let them know they can respond. During the first two practice trials, if participants provided responses before the videos ended, they were told "remember that to play the game you have to wait until the screen turns blue." Thus, the task itself was similar to the Video Only condition of Study 1: only the video, and no linguistic hints, were provided to participants.

The items were presented in blocks such that all the failures were presented one after the other, followed by all the successes or vice-versa. Participants in the Successes First condition watched all the success videos in one block, followed by a block of all the failure videos (with no clear break between blocks), while the Failures First condition did the reverse. This resulted in 22 videos described per participant.

Results

Negation

Each response was transcribed into text, and the resulting transcriptions were binarily coded according to the presence of a negator as in Experiments 1a-c. Children were much more likely to provide a negator in their response to failure scenes (64.5% of the time) as compared to success scenes (1.4%) in general (see Fig. 9). Participants who saw Successes First condition produce a few more negators (35.9%) than those in Failures First (30.45%).

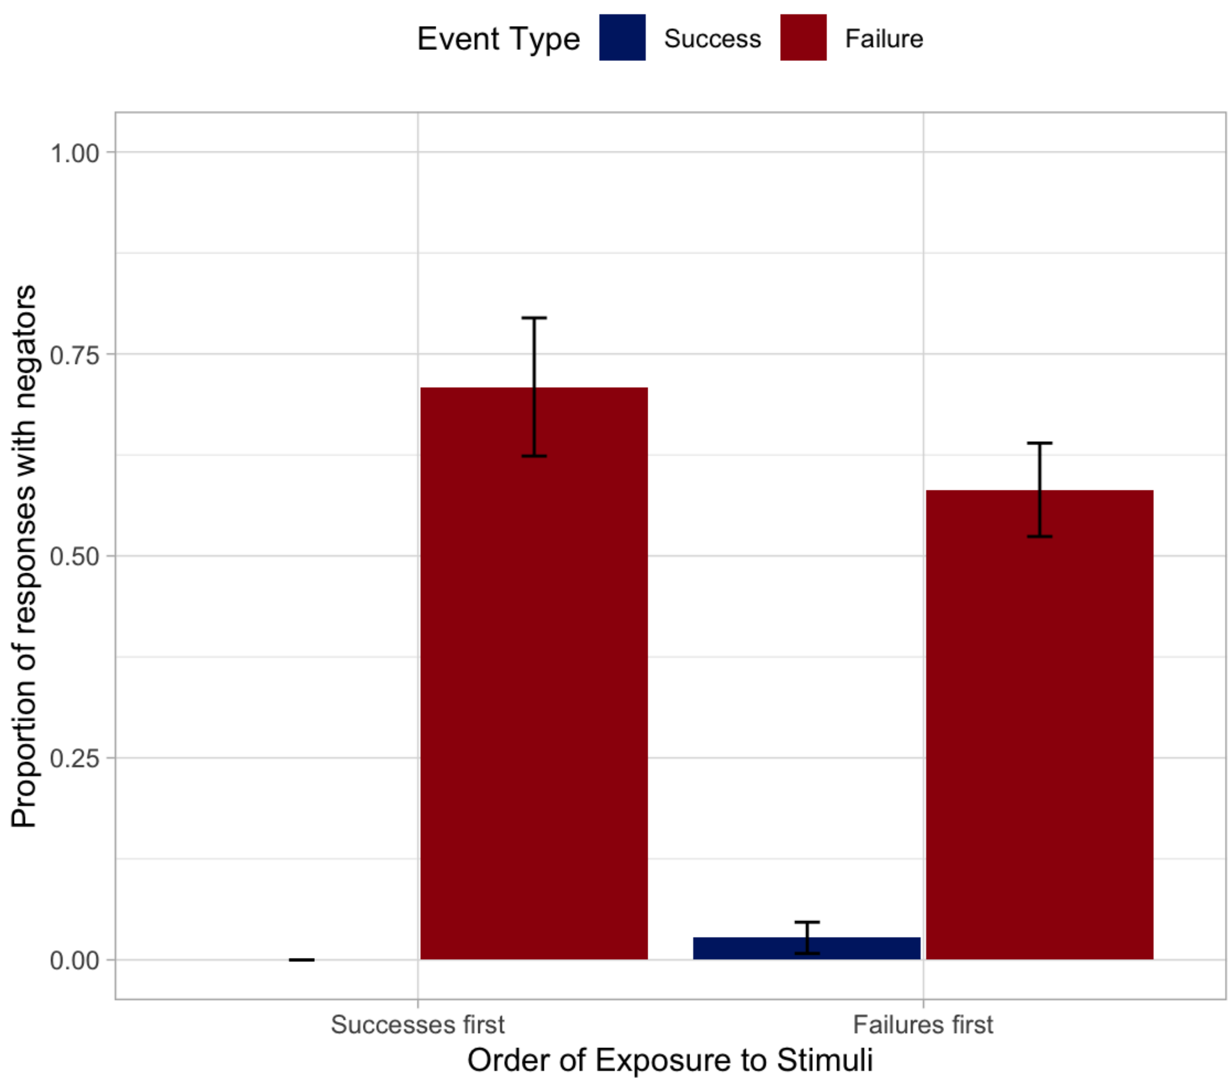

Figure 9: Proportion of child responses which included negators for Successes First and Failures First conditions. Study 2a. Data are presented by Event Type and Order of Exposure, with error bars indicating ±1 Standard Error.

As before, these observations were supported by a multilevel logistic regression using glmer to predict negator presence using Event Type and Block Order and their interaction as fixed effects, and by-participant and by-item random slopes as random effects. There was a main effect of Event Type as predicted, (β=-3.12, SE=0.609, z=-5.126, p<0.001) but did no effect of Order of Exposure (β=-0.208, SE=0.563, z=-0.369, p=0.712), nor an interaction between the two (β=-0.539, SE=0.426, z=-1.265, p=0.206).

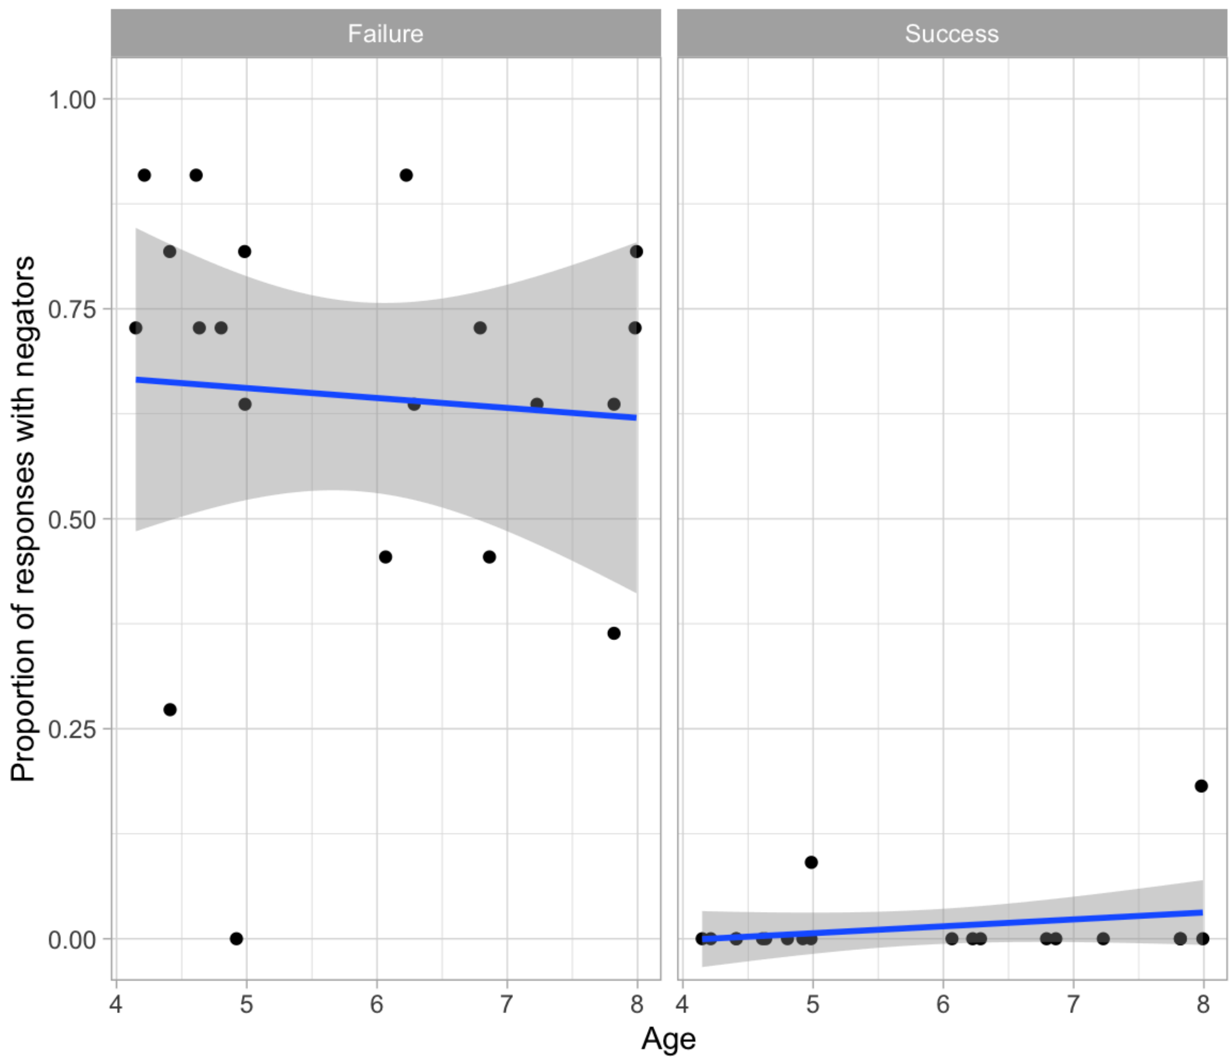

Figure 10: Proportion of participant responses which included negators for Failure and Success events by age. Study 2a. Data are presented by Event Type, with error bars indicating ±1 Standard Error.

Given the wide range of ages of the children participating in this study, we also conducted an analysis that included age in months as a predictor (see Fig. 10). This was not preregistered, and hence exploratory. While the participants in the Failures First condition were older (m_{age}=6.4 years) than those in the Successes First condition (m_{age}=5.3 years), these tests suggest that age is not a critical factor. A simple linear model using only age as a predictor found that it was not a significant predictor of negation responses produced ($\beta=-0.002$, $SE=0.0165$, $z=-0.110$, $p=0.913$). There were similarly no interactions between age and Event Type nor age and Order of Exposure.

Alternative constructions

The transcribed responses were binarily coded according to the presence of *try* as well as implicit negators (like *miss* or *fail*). The latter did not yield enough observations to merit analysis, with only one utterance including an implicit negator ("the girl was throwing it, but she kept on missing it again and again and again."). There were however a number of responses which remarked on the fact that someone was trying to complete an action. Generally, children produced *try* more often to describe Failures (44.29%) than Successes (3.27%), though many of these were not "alternative" descriptions in that they co-occurred with a negator. When considering responses in which there were no negators, children produced *try* more often in Failure videos (17.6%) than Success videos (3.3%).

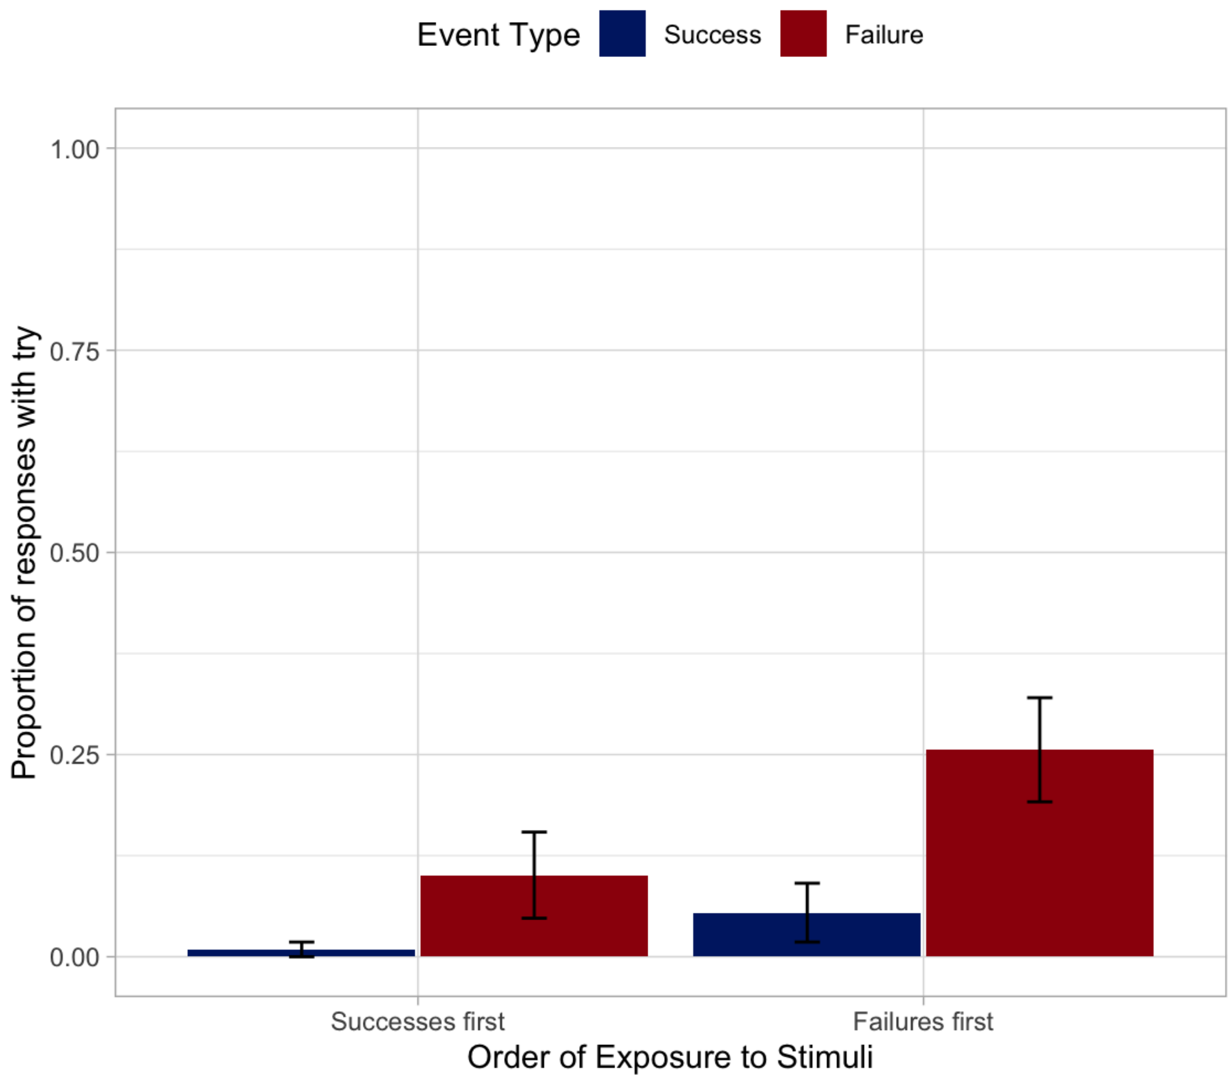

Figure 11: Proportion of child responses which included the word "try" but not a negator, for Successes First and Failures First conditions. Study 2a. Data are presented by Event Type and Order of Exposure, with error bars indicating ±1 Standard Error.

Some of these observations were supported by a model like that described in the previous section, but predicting *try* uses rather than negators. We found a main effect of Event Type (β=-17.101, SE=8.172, z=-2.093, $p<0.05$), but no effect of Order of Exposure (β=-2.916, SE=3.555, z=-0.82, p=0.412) nor an interaction (β=-2.284, SE=3.707, z=-0.616, p=0.538). We additionally conducted the same post-hoc analysis of age, and found that older participants were not significantly more likely to produce a description with *try* instead of using a negator

(β=0.015, SE=0.01, z=1.393, p=0.164), but they were more likely to produce *try* in general (β=0.034, SE=0.015, z=2.252, $p<0.05$).

Discussion

Using the same materials developed for Studies 1a-c, children were significantly more likely to freely describe a failure with a negator than they were a success. Though seeing an event successfully complete before seeing its failure seems to increase negator uses, this effect was not significant, suggesting the failure of simple actions may indeed serve as a null-discourse context for negator learning. Post-hoc analyses of age suggest that the older children that participated were not significantly more likely to negate than the younger children. Children were also significantly more likely to describe a failure with “try” than a success in cases where they did not negate. It is therefore also an available alternate construction for children attempting to describe failure which could likely be ruled out by syntax by more advanced learners (like adults in Studies 1b-c). Post-hoc analyses of effects of age found that older children were not more likely to say “try” instead of using a negator, but that they were significantly more likely to use “try” in general. This may be because, as discussed in the introduction of this section, “try” is more difficult than other verbs syntactically.

Study 2b: Adult Study

While children appear to use negators much more often than the adults in the Video Only conditions of Studies 1a-c, it’s important to remember that the child version of the task was greatly simplified. Given adult performance in Studies 1a-c, we predict adults in the child-friendly task will still regularly describe failures with negators, and that seeing Successes First will not significantly increase negator responses. We will statistically compare adult and child performance, as adults may be less likely to negate because they have greater inhibitory control (e.g., to inhibit the failed prediction) and a larger vocabulary (e.g., to generate alternative constructions with). For present purposes, no claim hinges on adults and children producing negators at the same rate in response to negators, but such comparisons may still serve informative in developing an understanding of failure as a null-discourse context for negator learning.

Methods

Participants

Twenty English-speaking adults participated and contributed data to the results reported below. All participants were recruited from Prolific and were offered $5 for their participation. Participants were excluded if they provided responses that indicate they did not understand the instructions (i.e., providing no response, or inaudible responses) more than 20% of the time. Two additional participants participated but were excluded due to persistent audio issues. Participants were randomly assigned to either the Failures First or Successes First condition.

Materials

The materials used were the same as those described in Study 2a.

Procedure

Participants completed a self-guided version of the same task described in Study 2a. After each video ended, and they spoke their response aloud, participants pressed the space key to advance to the next video. There was less control over the quality of the audio than in Study 2a, though this did not lead to issues in transcription. Other than these differences arising out of being conducted remotely, the only important distinction from Study 2a is that the adults were told that they would be completing a task designed for children, and that this task was to watch and describe a series of videos. This differs importantly from the instructions Studies 1a-c where participants were asked to guess how a parent would describe the videos.

Results

Negation

Responses were binarily coded for presence of negators as described in Study 2a. As Fig. 12 shows, adults were again more likely to produce negators to describe failures (55.45%) than successes (4.54%). And, again, participants produced more negators when exposed to Successes first (32.73%) than when exposed to Failures first (27.27%), though this difference is less stark.

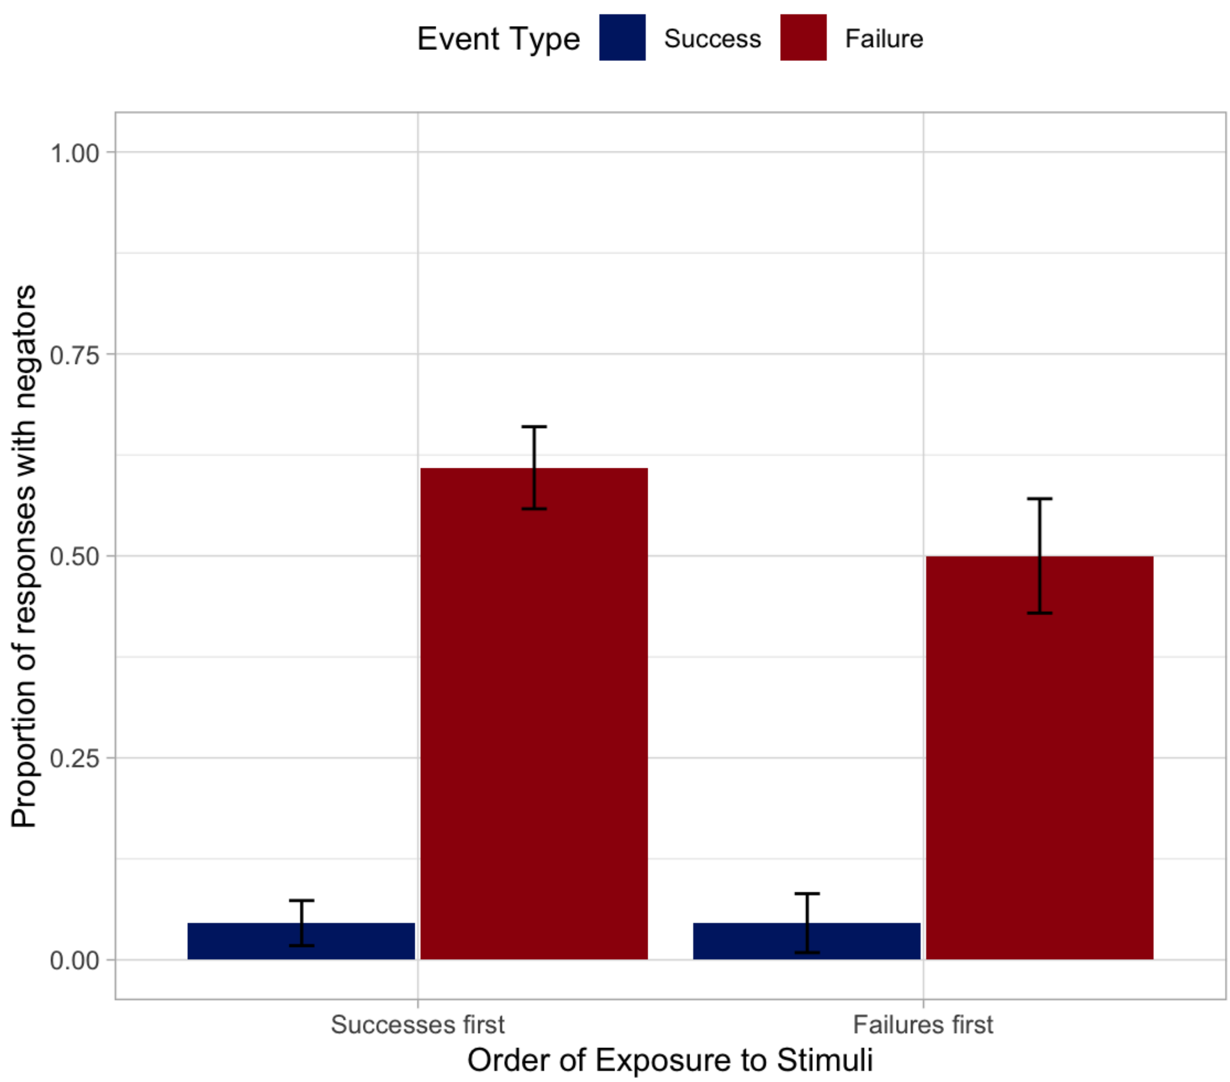

Figure 12: Proportion of adult responses which included negators for Successes First and Failures First conditions. Study 2b. Data are presented by Event Type and Order of Exposure, with error bars indicating ±1 Standard Error.

To investigate these observations, we used the same model described in Study 2a. We found a main effect of Event Type (β=-1.8199 , SE=0.244, z=-7.461, $p<0.001$), but again did not find an effect of Order of Exposure (β=0.147, SE=0.2653, z=0.553, p=0.58), nor an interaction (β=-0.131, SE=0.199, z=-0.659, p=0.51).

Alternative constructions

Like children, adults produced “try” more often to describe Failures (33.64%) than Successes (4.09%) in general. As Fig. 13 shows, they were also like children in that they were more likely to use “try” instead of a negator in cases of Failures (14.55%) than Successes (3.64%). While more uses of implicit negators were observed in adults than in children, the number remains too low to be useful for analysis (see Appendix D for descriptive statistics).

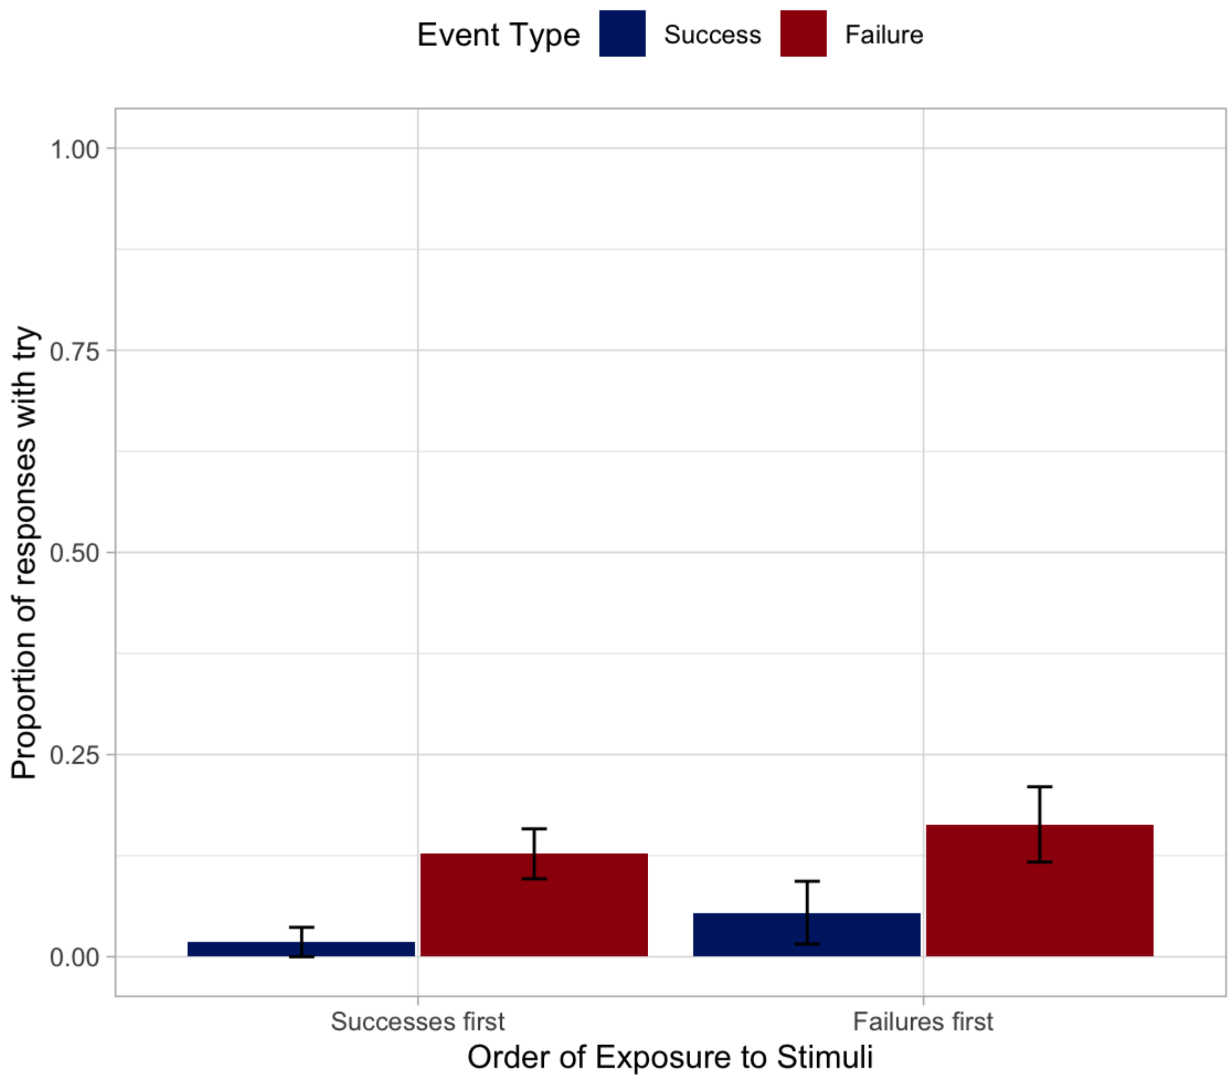

Figure 13: Proportion of adult responses which included "try" but not a negator for Successes First and Failures First conditions. Study 2b. Data are presented by Event Type and First Block, with error bars indicating ±1 Standard Error.

As for Study 2a, these observations were probed using a multilevel logistic regression to predict the use of "try" in the absence of a negator using Event Type and Order of Exposure and their interaction as fixed effects, and by-participant and by-item random slopes as random effects. We found a main effect of Event Type (β=-4.5999, SE=1.688, z=-2.726, $p<0.05$) but no effect of Order of Exposure (β=-0.282, SE=0.812, z=-0.347, p=0.729) nor an interaction (β=-0.083, SE=0.848, z=-0.098, p=0.922).

Developmental comparisons

Developmental comparisons were made using the same model described in Studies 2a-b, with age group added as an interaction term. This model found a main effect of Event Type ($\beta=-2.223$, $SE=0.204$, $z=-10.884$, $p<0.001$) and an interaction between Event Type and Age Group ($\beta=-0.3997$, $SE= 0.178$, $z=-2.240$, $p< 0.05$). In other words, children actually negated significantly more than adults did when presented with Failure Videos. Models comparing adult and child "try" productions failed to converge.

Discussion

Adults were again much more likely to describe failures with negators than successes, even when the videos they were asked to describe were presented in blocks and without fillers. Similarly, they were more likely to describe failures than successes with "try" when not using negation. Like children, and as Studies 1a-c would suggest, they also did not benefit significantly from viewing Successes First. In this simplified version of the task, participants appear to negate at higher rates than those in the Video Only conditions of Studies 1a-c, however analyses did reveal an interaction such that children actually negated more when describing failures than adults did. This is likely because adults have a larger lexicon and are therefore able to entertain a wider range of constructions in an open-ended production task with low cognitive load. In a number of the Failure trials, adults used "fail" (n=8) as well as morphological negators (i.e., *un-*, n=5) instead of *no* or *not.* However, such responses were not attested in child participants at all.

Study 2 Discussion

We found that, like adults, children also reliably describe failures in every day actions with negators in the absence of additional linguistic information. Unlike the typical methods for eliciting negators, our task took advantage of failure as a null-discourse context - one that is likely to occur in everyday contexts that adults were again shown to be likely to use negation in.

As items were separated into two blocks, this does not speak to contexts like those typically used in eliciting negation, in which a successful action immediately precedes a failure. It does, however, provide stronger evidence that young children are still readily able to recognize failed attempts to complete them despite presumably having less experience with these particular kinds of events than the adults in Studies 1a-c. This may be because they have already attained sufficient experience at the ages tested, or because the sorts of cues event perception relies on are more generally available at this age as recent work on event perception suggests. As the events chosen for the visual stimuli were simple, everyday actions, the former is a live possibility. In the strongest form, such a view would predict that younger children (e.g., infants, toddlers) would struggle with recognizing failures of these simple actions as such when they occur in isolation. It may also additionally predict that observing the successful completion of an action prior to its failure would be necessary for such learners. However, work outside event perception suggests that there is more going on than learning expectations, as even infants (14 months) are able to recognize when others cannot accomplish a goal (Warneken & Tomasello, 2009).

Finally, we found that children were significantly more likely to describe failures with negators in this task than adults. This could be because making the task child-friendly simplified it in a way that freed up our adult participant's already greater inhibitory control, and therefore made it easier to consider alternative descriptions (e.g., discussing what prevented someone from accomplishing something). However, the simplifications made some adult participants more aware of blocking and motivated more varied responses. Importantly, adults in Study 2b are negating more often than those in the Video Only conditions of Experiments 1a-c, confirming that the simplified version of the task succeeding in reducing other performance demands.

General Discussion

This work was designed to examine a specific hypothesis about how truth-functional negation may be contextually supported in a way that might be of use to children learning negation. Specifically, it sought to investigate whether failures of simple actions might support the use and inference of negation across development on their own, or whether additional linguistic information remained necessary. In two studies, we found that both children (4-7yo) and adults regularly used negators to describe failures but not successes. Surprisingly, across the Video Only conditions in Studies 1a-c, adults, who were simulating learners with no access to linguistic evidence, were significantly more likely to describe Failures with negators than Successes. This effect was replicated in Studies 2a-b with both children and adults in a child-friendly version of the task. When all participants in the Video Only condition across Study 1 are averaged together, there was a difference of 40.97% in negator responses between Success (3.24%) and Failure (44.21%) events. In Study 2, there was a 63.1% difference between Success (1.4%) and Failure (64.5%) for children, and a 50.9% difference for adults (55.45%) than successes (4.54%). Thus, the proportion of responses with negators is reasonably larger for Failure events than in comparison to Success events, even for children despite having fewer experiences with these events than adults. Both cases also exceed the general base rate of negation in descriptive corpus analyses of parental speech (7.6% of sentences included either *not* or *no* as reported in Gomes et al., 2023) as well as adult speech and written English (2.5% and 4.6% of sentences respectively based on Switchboard (n = 177,804) Wall Street Journal (n=49,208) and portions of Penn Tree Bank (Marcus et al., 1993)). While Study 1 demonstrated that linguistic information can be additionally helpful for adults, it was not necessary for encouraging the inference of negation. Study 1 additionally suggests that linguistic information is only helpful when relatively complete, as other experiments have suggested (Gomes et al., 2023). We therefore found support for the view of failure as a null-discourse context for negation, despite the fact that negation is so abstract and the disparity between children and adults in their experience with the events at hand.

What we can learn from failure

Negation is an abstract concept often thought to be far removed from perceptual experience. As such one might expect that the learning of terms that express negation would require access to more than just the perceptual circumstances of use (Feiman et al., 2017). While recent research has found that linguistic evidence can help (Gomes et al., 2023), the naturalistic nature of this work leaves unclear which situational contexts linguistic information is most helpful in. This question is especially pressing because young learners likely struggle to keep track of expectations of others over long periods of time, or to infer an expectation from the use of a negator, in addition to everything else they are trying to do as they learn a language. The present studies suggest that failures may indeed serve as null discourse negation contexts that 1) occur naturally and 2) require minimal (if any) non-local information (e.g., no need to track belief states of others, other information in discourse, and etc.), which may make them ideal “gems” for negator learning (Trueswell & Gleitman, 2020). These studies were motivated by recent research in event perception that found that children are sensitive to abstract event categories, and thus failures may set up local contrasts between the predicted category and the actual outcome.

The reliable effect of Failure events provides empirical experimental support for the claims made by acquisitionists that have studied negation’s forms and functions. These researchers have considered Failure negation as comparable to an event version of Nonexistence negation to explain its early use by children (Choi, 1988; Pea, 1980). However, these proposals did not include a specific mechanism by virtue of which Failures may be more easily observed in the “here and now” (Choi, 1988). Since negation can be used to comment on various sorts of unfulfilled expectations, ranging from the simple (e.g., absence) to the more complex (e.g., denying statements others may believe), there are surely other mechanisms that may lead to the development of expectations, some of which go unfulfilled. For example, researchers have suggested that there are "appropriateness conditions" for negation, such that for a speaker to utter "Helen isn't home" requires that they believe their interlocutor expected the corresponding affirmative ("Helen is home"), said it, or had implied it in what they had recently said (or did; Clark, 2019). As Horn (1989) points out, all this requires is that the corresponding affirmative is in the common ground or discourse, regardless of how it entered it. Thus, expectations may result from many cognitive processes that are not as close to perception as has been argued presently for failures, and so surely the causes can range from low-level (e.g., motion cues to some event) to higher-level mechanisms (e.g., inferences based on social roles and world knowledge). Relatedly, they may vary in how sensitive they are to context (e.g., is an expectation of catching formed when the catcher is an extremely clumsy friend?). Indeed, as mentioned in more detail in the Discussion section of Study 1, the elicitation tasks typically used for negation seem to depend in large part in setting up such a context by making some contrast more salient. Naturally, some of these sources of expectations are likely to be less accessible to infants and become more accessible throughout development, as they gain greater linguistic, conceptual and world knowledge.

Linguistic scaffolding

In Study 1, we found that one needs relatively complete linguistic knowledge to improve one's ability to infer the use of negation in the context of event failures, similar to Denials (Gomes et al., 2023). This might suggest that a learner needs at least a decent grasp of their language's syntax and exposure to utterances with known words. As a learner's vocabulary increases, the likelihood of them encountering utterances with unknown words decreases, and so they are more likely, upon encountering event failures, to hear utterances with known words and the requisite syntax. However the work on Denials has additionally suggested that children's relative delay in producing *not* vs. *no* is a result of increasing access to linguistic information (Gomes et al., 2023). The present evidence that linguistic information is not necessary for the inference of negation in the case of failures might seem to contradict this claim. That is, Denials, which are often communicated with *not*, are difficult to guess from instances of observation alone, but Failures, which are also often communicated with *not* are not difficult to guess from instances of observation alone. However, it's important to note that in the present studies, participants were asked to provide an entire sentence. Had the participants in Study 2 been asked to guess a single word, it's unlikely that they would guess a negator in the absence of additional syntactic information (they would likely guess a noun). As a result, this would result in repeated issues for learning *not* until the learner develops a more sophisticated understanding of their language.

Languages afford multiple ways to negate. One can negate by a free (*not*, *no*) or bound (*un-*) morpheme, and there are words which have been argued to implicit negate (e.g., *stop*). The framework provided here speaks primarily to negating by free morphemes, as the basis of the stimuli were events which are lexicalized in English. This raises interesting questions about the ways in which words thought to implicitly negate are acquired under our proposal, as the relationship between lexicalization and event perception requires further refinement. In other words, are the sorts of act typically labeled with *miss* perceived as an event in the same way that *hit* may be? If not, then this raises age-old concerns about analyticity and definition (i.e., could an individual have one without the other, and which is the primitive). If so, then our proposal requires further elaboration along the lines of e.g., temporal asymmetries between implicit negatives and their matched affirmatives (one must have been trying to HIT to MISS, GOING to STOP) and thus considering the referential timing of events and descriptions. We tentatively suggest that event perception attends only to abstract features and not to the particulars of the events, so that lexicalization is downstream from perception rather than exerting any top-down influence over perception.

Conclusion

Though negation could be used to comment on various sorts of unfulfilled expectations that are unrelated to event perception (e.g., absence), Failures are ideal learning instances for acquiring it. This is because Failure events bundle an expectation and it's unfulfillment in close temporal proximity and, as a result, do not require memory retrieval (e.g., where keys are typically kept), or engaging in more complex forms of reasoning (e.g., what mom is looking for after work). We showed that Failures do indeed elicit descriptions which include negators more often than Successes, and this is likely reflected in circumstances of use. This suggests that, for speakers with an understanding of events, there is information present in the "here and now" during event failures which could allow a language learner to map the forms negation takes in a given language to truth-functional negation. Despite its high degree of abstractness, negation, then, could be learned from observation, just as nouns and verbs may be (with additional nuance, of course). This may be why they are able to use it regularly in certain non-truth-functional contexts prior to age 2 and a half (e.g., absence, rejection, failure). As these are cases where the expectation is more transparent to an observer of a scene than, say, truth-functional negation which may require more word knowledge or sophisticated theory of mind abilities.

By focusing on failures, we argue that even a more limited learner, one who does not have extensive world knowledge or fully developed theory of mind abilities, would nevertheless find some instances of unfulfilled expectations that are commented on to be informative learning circumstances. Under other circumstances, however, such as those where parents use negation to describe situations may be the result of more complex expectations about the world (e.g., due to rarity or because it requires privileged information) those instances will be confusing, difficult to understand, and unlikely to serve as a learning instance for the child. In other words, the same considerations for content words apply for negation. Parents that "talk with" their child about the world around them are more likely to comment on a shared expectation that went unfulfilled, while those that talk “at their” child may discuss expectations that are not jointly-accessible in the “here and now.”

When learners are able to take advantage of the additional linguistic information provided by the syntax and semantics, the potential informativeness of Failure events is particularly high. Thus children may show variation in when they begin producing and comprehending *not*, not just due to their limited knowledge of syntax (Bloom, 1968) or differences in the circumstances they encountered during observation (Nordmeyer & Frank, 2018), but also because of individual differences in how their parents use *not* to talk about the world. If this were true, then it would be no surprise most children seem to have uncontroversially acquired *not* by 2-and-a-half years of age (Feiman et al., 2017; Austin et al., 2017) as this is when they tend to leave the two-word stage and reach a sophisticated understanding of the syntax of their language. At this point, a new source of information is available that can assist word learning in more referentially opaque contexts, providing each word learning instance with additional evidence to constrain the sets of hypotheses available.

CHAPTER 3: NEGATIONS ARE NOT CONTRARIES

THE PROBLEMS LOGICAL DEFINITIONS POSE TO WORD LEARNING

> *Son of Hipponicus, there is an ancient saying, that "hard is the knowledge of the good." And the knowledge of names is a great part of knowledge. If I had not been poor, I might have heard the fifty-drachma course of the great Prodicus, which is a complete education in grammar and language - these are his own words - and then I should have been at once able to answer your question about the correctness of names. But, indeed, I have only heard the single-drachma course, and therefore, I do not know the truth about such matters; I will, however, gladly assist you and Cratylus in the investigation of them.*
>
> - Socrates to Hermogenes in the Cratylus (Jowlett, 2019)

According to some accounts of creation, in the beginning, there was the word and the world followed shortly thereafter. Infants enter the world, but unlike Adam, infants cannot simply bestow the names of animals and expect others to follow suit. They must instead discover the words already in use by others around them. Infants learning their first language can therefore be said to experience creation backward. However, a lot has happened since the fall - and the world has drastically changed both physically and culturally. Could Adam think thoughts about carburetors or phlogiston? He never encountered the former, and the latter required a few paradigm shifts in science (and are a few paradigm shifts behind us). In cognitive science, this deep observation is often translated into questions about concepts, or how meaning is represented. Naturally, this line of research runs into deep questions about what meanings are in the first place,[6] particularly about nativism - whether we are born ready to represent what we encounter.[7] Unlike other animals, e.g., vervets (Seyfarth & Cheney, 1990), we are not limited to an innate stock of call-meaning pairs and our language reflects many of these changes. We can

---

[6] The question about meanings of general terms (i.e., labels that are not names) is intimately related to metaphysical debates about what sort of being thoughts have. Plato exemplified the realist position (forms (like The Good) have being) and similarly placed this being in a “third realm”(like Frege (1884)). Cognitive scientists, however, are not allowed into the third realm to check, so they typically follow in the footsteps of moderate realists (e.g., Aristotle (1975) or Quine (1955)) such views are easily compatible with materialist frameworks, as is e.g., conceptualism (Ockham, 1349).

[7] This does not imply it is present at birth, but includes maturational processes. Though particular inputs may be necessary at certain points (e.g., food for puberty), this does not involve learning.

and regularly do make up words and this undeniably makes them easier to use in reasoning (e.g., the engineer may be stalling because of the carburetor). Even if one allows for an innate set of concepts, it could not possibly include meanings that refer to things created after the fall.

Acquisitionists, not being mechanics, focus on words other than carburetor, like one or not. This is because these words have been formalized in certain theories (e.g., math, logic) which allows for greater precision than category words (or sortals, e.g., dog). This was, relatedly, the result of various paradigm shifts in how we reason (e.g., the computational revolution; Johnson-Laird, 1999) and so concept learning theorists argue learning such words may be good candidates for identifying conceptual discontinuities (Carey, 2009). Infants do seem to learn some kinds of words earlier than others (e.g., nouns before verbs though see Gillette et al., 1999), and this has been argued as suggestive that learning some words requires more than just mapping it onto a meaning (Gentner, 1981; Carey & Levine, 1982). But the formal words specifically exhibit interesting learning trajectories, where some words (e.g., one, no) come before others (three, not). This linguistic work has in turn further motivated non-linguistic work to probe the child's pre-linguistic concepts (Carey, 2009 for review on number). Though alternative conceptual continuity hypotheses have been proposed for the linguistic data (Feiman et al., 2017; Gomes et al., 2023), they rarely speak to the non-linguistic findings in these domains. The goal of this paper is to review both sets of developmental patterns for negation (linguistic and conceptual) and to suggest an alternative that assumes conceptual continuity.

The alternative explanation for word learning patterns of formal words is naturally to argue that they are just that, word-learning patterns. It has perhaps been most clearly articulated as the result of a linguistic limit (Feiman et al., 2017) where learning different parts of language learning results in an "information change" which makes a host of previously available but unlikely hypotheses more easily accessible (Gillette et al., 1999). For example, knowing a bit about syntax can make it more likely that you will guess a verb is being said than a noun

(Gillette et al., 1999; for a review, see Gleitman et al., 2005). In essence, it suggests that the different levels of linguistic organization provide different kinds of information that the learner can exploit. For example, morphological and syntactic information can help infants learn aspects of word meaning, which has been extensively studied for verbs (e.g., transitivity; Yuan & Fisher, 2009). Some word classes (like verbs) may therefore arise later because they benefit from different kinds of information (syntax), not because they are more complex conceptually. Therefore, only a learner who has already learned some aspects of their language's syntax will have access to this information for further word learning. This was initially articulated and demonstrated for verbs (Gillette et al., 1999) and a similar experiment similarly supports an Information Change account for the linguistic data.

Because Information Change is an explanation of word learning specifically, it cannot explain developmental changes observed in non-linguistic tasks. However, Information Change does suggest a broader explanation that could account for developmental changes observed in both non-linguistic tasks and linguistic tasks. Since Information Change proponents construe the referential world as rife with ambiguity, they repeatedly frame syntax as a "zoom lens" which helps bring meanings into focus. This does not suggest that syntax is meaning-constitutive for verbs,[8] but rather that syntax helps your interlocutors access some construal in the moment. For example, an infant observing a thinker engaged in active thought may not be compelled to access THINK thoughts unless they have reason to think e.g., a predicate is under discussion and that it allows for sentential complements. Similarly, an infant seeing an whale may not be compelled to access MAMMAL thoughts unless they already know a whale is a kind of mammal. I will argue that this view, access discontinuity, suggests that the conditions under which we access a concept can, and in fact must, change through the process of word learning. Over-analogizing to similar questions in syntax has resulted in a focus on the performance/

[8] Though this is, of course, not to deny the possibility of a deep relationship between syntax and semantics.

competence distinction, but what constitutes competence is what is under debate in the first case. The conceptual literature is not settled on whether competence reduces to access. While some theories of concepts have defined conceptual competence this way (e.g., Locke, 1690), others do not (Fodor, 1975; 1987; 1988; Fodor et al., 1981). Without equating access and competence, someone failing to access a concept in some context does not serve as knock-down evidence that this individual is incapable of representing it. This is not just a matter of being able to represent something, but of accessing it in that moment. I will argue that this assumption is not warranted, especially for words that are not straightforward categories and instead operate at the “message level.” The non-linguistic tasks are thus inconclusive, as it requires assuming that the trials were constructed such that if one has the given concept (e.g., negation) then one would access it in those moments.

We will begin with an introduction to negation: its logical status, developmental data (production and comprehension), and attempts to probe its presence non-linguistically. We will then introduce the concept literature and concept learning proposals from different theoretical backgrounds and the notion of protonegation. I will briefly cover the standard arguments against concept learning, namely that it is not amenable to the only model of learning at hand: hypothesis testing. I will also briefly argue that rather than integrating and allowing space for social construction, approaches that reduce or otherwise identify concept access with competence risk “semanticizing” social constructs (e.g., naturalizing a definition via semantic/cognitive claims). I will then argue that existing findings can be explained by Access Discontinuity, which is more amenable to modeling as hypothesis testing, and less liable to runaway naturalization. Access Discontinuity leaves the board such that in addition to addressing these limitations, concept learning theorists need to either argue 1) that competence totally reduces to access, or, more conservatively, 2) that it does in the experiments in question.

Conceptual Concerns

Most discussions on concepts begin by defining the Classical Theory of concepts, and noting where other theories deviate as a response to issues with either the empirical or philosophical commitments of the theory and I do not see any reason to buck this trend. Though there is no singular theorist or theory that is captured by the Classical Theory, it is highly associated with the British Empiricism and went on to influence various other philosophical and psychological traditions. At its most general, The Classical Theory simply states that 1) there are ideas (i.e., concepts for our purposes) some of which are simple and others complex, 2) that the latter must be constructed from the former, and 3) the simple ideas are somehow given (i.e., available through some non-learning process, e.g., innate). Most word meanings in a natural language are argued to correspond to complex ideas, and can be construed as being defined by other complex or simple ideas but must remain ultimately reducible to simple ideas. The process of constructing complex ideas may properly be called learning, since it involves the deployment of reason and is often construed as deliberate. Since complex ideas are learned by combining simple ideas, and simple ideas are themselves not so structured (or they would be complex ideas), it follows that they are not learned. Therefore simple ideas are often taken to be developmentally primitive, or in other words, innate concepts. Because of their simplicity and primacy, simple ideas are more readily accessible than complex concepts both to learning and thought (e.g., DOG comes to mind as a description for a DOG before ANIMAL). The mechanism for delineating simple ideas from complex ones was largely based on intuitions about what information is available to the senses (or to perception, as later theorists would claim). For example, RED is a simple idea as it is the output of some set of sensations impinging on the retina, but PLANET is not. Upon seeing Mars in the night sky, the classical theory predicts we would think the former more immediately than the latter, though this could be attenuated with practice.

What makes the Classical Theory so appealing is that it provides elegant solutions to the problems of extension (how do concepts refer to things in the world; because there is a sensory bedrock underlying them), categorization (how do we determine that an item is of a category; it has the necessary and sufficient features given in the category's conceptual definition), and concept acquisition (how do children learn new words; they construct definitions for them from primitive features). However, despite these virtues various concerns have been raised, not the least of which is that no one has successfully offered a definition of a complex concept (Margolis & Laurence, 1999), and it's not for lack of trying. Every year, legions of undergraduates throughout the world take introductory philosophy courses where they are tasked with providing a set of criteria that includes chairs and excludes non-chairs, and every year the instructors respond with examples of chairs that are excluded or non-chairs that are included by the criteria. This is not due to the student's lack of experience, as even tenured philosophers with book deals fail to provide a compelling definition of chair (Elbourne, 2011). The problem is not any easier for verbs, as Fodor (1983) illustrated with paint. Could the meaning of paint be to cover with paint? What about a horrible accident at the paint factory that results in the spectators being covered in the pantone color of the year? Could paint mean to cover with paint intentionally (i.e., some agent intends to cover some surface with paint)? Many would bristle at the idea that Michaelangelo simply painted the ceiling of the Sistine Chapel, he was painting a picture on the ceiling. But even then, what about when he dipped his paintbrush into the paint? Surely he intended to cover it with paint, but it would be very odd to claim he painted it in so doing. In addition to difficulties providing a definition that does not easily fall prey to counterexamples, the psychological reality of definitional structure has been brought under scrutiny after failures to find differences in processing between ideas of differing complexity (Kintsch, 1974). In other words, BACHELOR should take longer to process than MAN since there must be a higher number of necessary and sufficient conditions for its deployment, but it does not (Fodor, Garrett, Walker & Parkes 1980). These concerns, however, were compounded with the death of the analytic-synthetic distinction.

Briefly put, its demise went like this: a statement may be true for one of two reasons, due to the nature of the terms themselves (analytic), or because they say something true about the world (synthetic). For example, you do not need to look to the world to determine whether someone being a bachelor makes them an unmarried man, but you do to determine whether a particular individual is a bachelor (e.g., by asking them or others whether they are married). Quine (1951) famously criticized this distinction by arguing that analyticity rests on an assumption of synonymy between a term and its definition such that they are interchangeable (e.g., BACHELOR could be subbed in with UNMARRIED MAN in any context, and vice-versa). But, to determine whether a term and definition are interchangeable requires a notion of necessity to distinguish accidental coextension from those resulting from the meanings (Quine, 1951). For example, in “Necessarily all and only creatures with a heart are creatures with kidneys,” both creatures have the same extension, but no one would argue this is an analytic fact (unlike “Necessarily all and only bachelors are unmarried men.”). To Quine this meant there was a vicious circularity in the distinction: analyticity requires synonymy, synonymy requires interchangeability, but how is interchangeability itself determined? With analyticity gone, so with it goes a straight-forward distinction between matters of meaning and matters of experience, and the ground it set for introducing the arguments for Mad Dog Nativism.

Mad Dog Nativism

Quine’s argument against the analytic/synthetic distinction included many arguments for meaning holism, or the idea that a term’s meaning is a consequence of a network of relations to other terms. This may be all well and good for describing science, but Fodor worried that holism would make language learning largely intractable (and perhaps, additionally, fails to provide an explanation over a description). Fodor’s desire to ensure that the possession of a primitive concept is not contingent on the possession of another concept led him to argue that words correspond to atomic (simple) concepts, instead of structured ones (be the structure definitional, featural, or holistic). Fodor's argument (for "Mad Dog Nativism"), though often seen as outrageous, runs simply. As characterized by Carey (2014), it runs:

> P1: (Confirmation) All learning is hypothesis confirmation.
> P2: (Logical Construction) One can learn new concepts only by creating and confirming hypotheses formulated in terms of logical constructions from antecedently available primitive concepts.
> P3: (Atomism) Most mono-morphemes cannot be analyzed as logical constructions of other concepts primitive or otherwise…
> Conclusion: (Innateness) In order to acquire a new mono-morphemic concept, one would have to confirm hypotheses already containing the mono-morpheme (i.e. already containing the concept to be learnt).
> Therefore, most mono-morphemic concepts cannot be learnt.

Fodor goes to pains to clarify that this is not to say that there are no complex concepts at all, but simply that they do not map onto single morphemes. Complex concepts are in fact contingent on the possession of primitive concepts, but that's all well and good, and achieved by the same sort of productivity that allows for natural language (e.g., limited set of words that compose in ways indicated by syntax). BROWN COW is a complex concept, and so is DRINK MORE OVALTINE. If these propositions could map on to words, then Fodor worries that the floodgates are opened and every bit of data needs to be considered during confirmation (meaning holism) and we are left with an intractable learning problem that also seems at odds with how meanings work. In the case of the latter, you and I can disagree about the definitions or features or permissible inferences of BACHELOR (is the pope a bachelor?) while knowing it's BACHELOR-hood we are disagreeing about). In the case of the former, it does at one extreme by limited any potential for the disjunction problem: that sometimes concepts (e.g., COW) are triggered in error (e.g., by a horse), and thus should perhaps cow can be understood as meaning COW-OR-HORSE as they can be tokened by cows or horses (Fodor, 1989). To give an example more relevant to word learning, learning bat would be a lot harder if they could learn not just that it could mean BATANIMAL or BATWEAPON, but also BATANIMAL-OR- BATWEAPON, and potentially even BATANIMAL-OR-BATWEAPON-OR-BIRD-AT-NIGHT. Or, for superordinate terms, learners may consider hypotheses that involve disjuncts of all the categories they include (e.g., MAMMAL = BAT-AND-DOG-AND-CAT- AND-WHALE…) or permutations thereof (e.g., MAMMAL = DOG-AND-CAT…).

Fodor's concerns about holism are especially clear in his debates with the Churchlands on meaning (Churchland, P.S., 1978; Churchland, P.M., 1981, 1988; Fodor, 1988; Fodor & Lepore, 1999) as well as just about everything else he's written. Indeed, this aversion to holism is in part motivation for the adoption of Informational Semantics as a tenet of his form of nativism. Holism, particularly as conceptual role semantics, allow for the use of words to figure in their meaning, but Fodor feels this conflates semantics with epistemology. For example, though we know dogs bark and as such bark and dog are used in similar contexts does not mean that this relationship should be constitutive of what DOG and BARK mean, only what we happen to know about them.[9] Informational Semantics on the other hand, solely posits that DOG means dog is a result of nomic (lawful) connections between their doghood (their being dogs) and "being causes of actual and possible DOG tokenings in us (Fodor, 1998)." This focus on nomic connections means that when learning monomorphemic words, children need not consider relations with other words (i.e., the existence of one need not be contingent on the existence of another). In a sense, the learning of all mono-morphemic words could be described as merely a problem of mapping labels onto existing concepts, instead of constructing concepts with the assistance of labels. With the argument for Mad-Dog Nativism in place, we are now more equipped to discuss what exactly a conceptual discontinuity might be and develop our theory of access.

The notion that all words are concepts and most word learning is just a mapping problem is a tough pill to swallow. It does not do anything to limit the initial conceptual state, and so many argue it is inherently non-parsimonious (Carey, 2009; Cowie, 1997). And, as many have

[9] As a commenter pointed out, this is true even if every time an individual tokens DOG they also token BARK because they have built up a strong association between barking and dogs.

pointed out, it potentially makes picking out the right concept harder if, along with good relevant concepts (e.g., QUANTA) that have scientific explanatory adequacy, we get bad ones (e.g., PHLOGISTON; Carey, 2009). Additionally, some have taken pains to argue that P2 (Logical Construction) is unfounded (Carey, 2015) and suggested alternative models for learning atomic concepts (Carey, 2009). This last criticism is especially relevant to our current question, as P2 seems to assume some sort of logical machinery (which presumably represents things like negation).

Mad Dog Nativism and Logic

When Fodor has presented his arguments for Mad Dog Nativism he repeatedly took great care to specify that most monomorphemes cannot be given a logical analysis when stating some variation of P3 above (Atomism), and, relatedly, when stating concluding that most monomorphemic concepts are innate. Presumably he specifies most rather than all in large part because of monomorphemic logical concepts, those which are critical for P2 which he grants can be reasonably defined. In Fodor (1992), he seems to accept a definition for logical concepts rooted in its inferential role: “Correspondingly , it is arguably a sufficient condition for a speaker' s meaning… ‘and’ that , ceteris paribus, he takes ‘P and Q’ to be true iff he takes ‘P’ to be true and ‘Q’ to be true; and so forth.” Fodor (2004) is more explicit on this point: “Since these [logical constants] are, plausibly, not referring expressions, it might be that an account of the rules that determine their conceptual roles is the whole story about why they mean what they do and what it is to understand them.” However, Fodor ultimately argues that one should not mistake a theory of what a constant is (e.g., AND, NOT) for a theory of how it is “learned or what it is to understand it.” It would therefore be a mistake to read this as a weakening of his nativist position, especially in light of earlier works, like “On the Impossibility of Acquiring ‘More Powerful’ Structures” which explicitly argues that it is impossible to learn a richer logic from a weaker logic (Fodor, 1980).

As the quotes in the previous paragraph suggest, it is common to consider the meaning of a logical constant (i.e., its inferential role) as a consequence of its Introduction and Elimination Rules. These “characteristic implications” are similarly taken to be the meanings in natural deduction approaches (Aspeitia, 2007). Unlike the previous examples, however, these rules can be constructed so as to avoid obvious circularity. For example, Peacocke (1996) suggested a conceptual framework wherein the meaning of and (AND) is defined as “that concept C to possess which a thinker must find transitions that are instances of the following forms primitively compelling, and must do so because they are” are of the form displayed in Table 3.

Premise	p		
Premise	q	pCq	pCq
Conclusion	pCq	p	q

Table 3: A table demonstrating elimination and introduction rules for conjunction. From two separate premises, one may either introduce a conjunction (as in column 1), and from a conjunctive premise, they can eliminate conjunction to conclude either premise alone.

That is, such a thinker that is capable of thinking about AND must find these introduction (i.e., from thinking p, thinking q to thinking pCq) and elimination (i.e., from thinking pCq, they must be able to think either just p or just q) rules for conjunction compelling. This is why Fodor argues this approach attempts to reduce meaning to epistemic or dispositional states: possessing C is being disposed to thinking solely p (or q) from pCq and vice-versa. Before reviewing his more general critique, it’s relevant to note that such approaches run into specific implementation problems. As Harman (1986a) demonstrated, it is easier to give a natural deduction account for AND than other logical constants. For example, while it is fairly straightforward to state Introduction Rules for OR (p implies OR(p,q), q implies OR(p,q)), the Elimination Rules are more complex since disjunction does not imply the truth of particular disjuncts, only that at least one will hold (OR(p,q) does not imply p and it does not imply q, it implies p or q). Thus, he suggests a sequent calculus approach rather than natural deduction. In such a logic, an argument is valid iff its premises cannot be true without one of its conclusions being true, and OR can be defined (OR(p,q) ⊢ p, q), that is its conclusions are understood disjunctively. A natural deduction account of NOT similarly runs into issues but with tractability (if p, NOT(p), imply anything; if p & other premises imply everything, then premises imply NOT(p)), and thus requires further "extend[ing] the notion of implication in various seemingly ad hoc ways."

Despite suggesting a system that can generate relevant truth-functional responses, Harman is unsatisfied because such an approach does not distinguish between logically equivalent connectives (AND(p,q) = NOT(OR,(NOT(p) NOT(q))) because they have the same characteristic implications (p, q). However, he suggests such connectives would remain different in meaning as a consequence of how "immediate" the implications are (e.g., AND(p,q) immediately implies p,q while NOT(OR,(NOT(p) NOT(q))) immediately implies OR(NOT(p),NOT(q)) presumably). Immediate implication is difficult to sort out formally, in part, Harman argues, because it "is not just a matter of the contribution a construction makes to truth conditions," but rather a matter of psychology, and thus has to do with an individual's dispositions. As Harman (1986b) notes, one cannot identify simply immediate implication (1) with disposition (2). This may be for simple reasons where an immediate implication is overridden (e.g., by a more immediate one, or performance factors), but also cases where an immediate implication is absurd.

1. P, Q, …, R immediately imply S for an Agent.
2. An Agent is disposed to treat P, Q, …, R as a reason to believe S.

Harman does not seem to imply this is an intractable or otherwise insurmountable problem, and instead concludes it means that truth-conditional considerations do not matter more for logical than non-logical words.[10] However, this is exactly what Fodor (2004) is objecting to when he

---

[10] That is, Harman thinks the meanings of both logical and non-logical words have meaning by virtue of their Conceptual Role.

argues that concept possession cannot be reduced to dispositional or epistemic states. We turn to this criticism now.

Peacocke's (1996) definition of AND requires that the thinker find the characteristic implications "primitively compelling" and can thus be likened to immediate implication (Harman, 1986a; Harman 1986b). The definition goes on to specify that they are recognized in virtue of their form, but Fodor (2004) argued it's unclear what form this is. It cannot be recognized in virtue of logical form as that would reintroduce the potential for circularity nor syntactic form as truth and validity are not syntactic notions. This would mean, as Harman suggests, the formal considerations are not all that matter for defining logical constants. If notions of immediate implication or primitive compulsion are contingent on dispositions or epistemic state, the potential for holism is reintroduced. It is additionally difficult to provide an account of immediacy that narrowly applies to logical words and not to non-logical words (e.g., prepositions like below which can be defined in relation to other prepositions like above) assuming one (like an early learner) does not know which category a novel word belongs to (Fodor & Lepore, 2006). Thus, the same arguments that lead to Mad-Dog Nativism about most concepts seems to apply to logical concepts despite the fact that they appear definable. This is because logical meaning relies on a notion of (truth) identity such that any operations which affect the same changes to the truth conditions of a sentence "mean the same thing," but, at the same time, it must explain why upon seeing P we think P first and not not not P. One way to do this is to suggest P is more primitive than not(not(p)),[11] but as Harman demonstrated, this is a difficult task and it seems unlikely that a purely formal account can be given. More practically, the classical story, the senses provide the grounds for identification, so how would such a learner identify the introduction and elimination rules that guarantee the meaning of AND? If monomorphemes can indeed map onto non-atomic concepts, then this likely introduces more problems as not only are

[11] Just as for classical theorists RABBIT might be more primitive than UNDETACHED-RABBIT-PARTS.

the many reasons (including non-semantic ones) one said AND, there are many ways of formalizing AND (Fodor & Lepore, 2006).

Formal definitions likely do not feel like enough because they do not unequivocally suggest what causes thinking of them. Unlike the definitions for categorical concepts like BACHELOR in the classical account, the definitions of logical constants are difficult to translate into a causal story (e.g., unmarried men (and whatever primitives they bottom out into) cause BACHELOR thoughts) and thus harder to envision word learning contexts for (e.g., utter bachelor around unmarried men). Fodor suggests that even if and can be defined according to its inferential role, that does not guarantee that its inferential role determines its possession conditions. Where inferential role is constrained by considerations of validity, the concept one possesses (at the moment) is constrained by one's mental state (at the moment). In other words, what makes something true is not necessarily what makes something a good instance of some category. This is perhaps clearer for non-logical words, where the latter but not the former is affected by e.g., visibility conditions since a cow in thick fog is still a cow despite being a poor instance of one.[12] Relatedly, one may be indisposed to accept an implication due to performance limitations like memory constraints, for example if the minor premise of a modus ponens is 5,000 words long. As we will see in Section 3, developmental psychologists take great care to account for the degree to which facts about performance may contribute to failure or success. Like the case of object permanence, what looks like a discontinuity in competence may be the result of discontinuities in performance (Baillargeon, 1985; Onishi & Baillargeon, 2005). By appealing to discontinuities in performance one can avoid committing to discontinuities in competence. This is because there may be no learning required (e.g., just maturation) depending on the nature of the mechanisms involved in performance. For the present conversation involving logical words, analogizing to the competence/performance distinction from syntax may obscure the real disagreements. Importantly, one needs to consider

[12] And why a horse in thick fog is never a cow despite being confusable for a poor instance of one.

the possibility that there is a discontinuity in how terms are accessed (i.e., Access Discontinuity), and this discontinuity may still involve learning - even if not of concepts themselves, instead the range of contexts others in the community tend to access them in.

What sort of discontinuity are we talkin'?

As a quick summary, we began this section by introducing the classical theory of concepts which states that there are some innate concepts and a hypothesis testing framework that can be used to define new complex concepts and thus guarantee truths of meaning. These innate concepts were sensory, thereby establishing reference, and thus truths of experience. However, the conception of definition at play is dependent on synonymy that is itself rooted in experience (e.g., any utterance of the definiens can replace any utterance of the definiendum and vice-versa). This presents problems to a first language learner, as they lack the meta-language necessary to avoid circularity. While it seemed that logical words might be exempt from the motivations behind Mad-Dog Nativism because their meanings are definable, their definitions rest on identity and this runs into similar issues that the classical theory runs into with synonymy. This in turn makes it difficult to reduce logical words to non-intentional terms without invoking notions that would re-introduce holism like "immediate implication" or "primitive compulsion."

The previous section ("Mad Dog Nativism and Logic") ended by suggesting that possession conditions are overdetermined, and thus discontinuities in possession may be a consequence of discontinuities in performance rather than competence. However, the performance/competence elides an important disagreement about what it means to have a concept, and as such has been poorly received. In a review of Carey (2009), Rey (2014) argued that there is a distinction between concept manifestation and concept possession, and defined them by analogy to performance and competence (respectively). Carey (2009) can be said to speak to manifestation, which Rey argues was not Fodor's focus, which he argues was concept possession. He argues that the latter is a prerequisite for the former: one must possess a concept before it is manifest, or “actually activated.” Carey (2014), in response, did not so much disagree with Rey’s distinction between concept possession and manifestation as much as see possession as “unproblematic, but empty.” This is because, while for Rey, the contexts in which one uses a concept falls under performance (or manifestation), for Carey, it is an important part of what it means to become able to use (or possess) some concepts. However, this way of putting the situation overcomplicates the terminology. First, Fodor has often said that concept individuation (i.e., what DOG is) is prior to concept possession (i.e., whatever it is that having the concept DOG consists in having; see Fodor, 1989). It would thus be more apt to analogize competence to concept individuation, especially given Rey’s characterization of possession as concerning the “logico-semantic issues of expressive power.”

At a basic level a claim of conceptual discontinuity is a claim at a discontinuity of individuation rather than possession, and so this convention will be adopted going forward. Second, analogizing to syntax might lead one to believe that if one disagrees with individuation discontinuity stances, one must thus take a performance discontinuity stance. Given the sorts of questions asked in syntactic development, performance typically refers to other qualities about the individual that may impact the execution of some competency (e.g., limited memory, poor motor skills). However, too narrowly construed, this makes it as seem as though the only live alternative to an individuation discontinuity account is a performance discontinuity account, and thus anyone who argues against the former must account for the evidence by recourse to the latter. That is, given some experimental data and interpretation that argues in favor of individuation discontinuity, the only course of action a proponent of individuation has is to argue that competency was masked by performance constraints (i.e., limited memory, inhibitory control ability) in the given task. In response, proponents of individuation discontinuity fairly point out that such arguments are not very useful if they do not suggest specific mechanisms (see p. 453 in Carey, 2009). While this is a fair rebuttal to vague criticisms of performance, cashing out the debate in this manner presents a false dilemma between individuation and performance discontinuity. This is because, as the last chapter suggested, however, there may be some non-performance explanation rooted in immediacy. We will now briefly introduce this potential account, dubbed Access Discontinuity.

While Harman (1986a; 1986b) proposed a simple formalization of immediacy to develop it as a notion, he ultimately concludes it must be a psychological notion rather than a formal one (as Peacocke (1996)). This is because, as Fodor (1986) suggests, these are in essence accounts of use, and use can be affected by a potentially open-ended number of factors (some cognitive, others not). Performance factors are only one subset of such factor, and they are likely of note at least in part for 1) theoretical reasons since (likely) many performance accounts are generalizable[13] and 2) methodological reasons as they can sometimes or in part be controlled for. Factors other than performance can affect whether a disposition actually holds (e.g., how the world happens to be at the moment), and so many are likely not amenable to psych study, however, immediacy, or access, is to some degree. Immediacy can be more simply understood as those sorts of situations in which speakers of a language typically access a concept. The study of psychology is, in large part, a study of access,[14] and as history shows it is difficult to develop a psychological theory of identity by manipulating access. There may be some cases where a concept is not accessed for performance reasons (as in the case of the 5000 conjunction modus ponens or during an interference task), but I will argue that importantly there are also cases that might result from accessing certain beliefs about one's community (e.g., for English speakers in Philadelphia in particular, thinking BACK DOOR likely causes them to immediately access STEP DOWN thoughts, or in the case of logical words generally having a label for it could help learners realize new contexts where it may be useful to access). While one could hold performance discontinuity views about some concept and believe learning is not

[13] Alcohol tends to affect linguistic behavior similarly across individuals, and drowsiness for similar reasons.

[14] Mostly because it's the one of the main things experimenters can control via the experimental environment.

implicated in the transition between stages (i.e., it is a result of typical maturation), this is not the case for an access discontinuity account. A strong Access Discontinuity account may, for example, involve standard associative learning (e.g., associating a tone with food) as long as it does not otherwise affect the underlying competence (e.g., whether/how a rat individuates that tone from other tones). That is, it primarily posits that the manner by which meanings are accessed are not to be identified with meanings themselves (see Halberda, 2019 for the same point in perception). While a strong Access Discontinuity stance is controversial, the weaker stance is commonplace. The simplest example may be the case of early word-learning itself, where children go from accessing DOG because they saw a dog to accessing DOG because they heard someone say "dog." Most importantly, the Access/Individuation distinction gets to the heart of the active sites of disagreement in the conceptual literature compared to the Performance/Competence (or Possession/Manifestation) distinction. This, of course, not to suggest performance factors play no role at, but that the philosophical discussion has not focused on it as psychologists do.

Thus far, it seems that to explain a behavioral discontinuity in linguistic development, the debate has been often presented as thought it must either be reduced to discontinuities in individuation or in performance. I argued there is another option: that how humans access concepts can change over time without necessitating any broader conceptual change. With the conceptual background in place, we can now turn to negation and ask whether present empirical evidence supports Individuation Discontinuity in light of recent non-linguistic research or whether an Access Discontinuity account is sufficient.

Negative Concerns

Though the most comprehensive concept learning proposal focuses largely on number words (Carey, 2009), this paper will focus on another abstract concept that's perhaps even more vexing: logical negation and corresponding negators in natural languages (e.g., *no* and *not* in

English). This is because, like number, negation is abstract (not say an object or event) and can be formalized in various systems, but unlike number words, negators are 1) perhaps even more basic to reasoning, 2) learned much earlier, and 3) not tied to any easily identifiable routines (like counting).

For advocates of concept learning, a standard research program is to 1) describe the “initial state” (what children can represent/process), 2) describe the adult state, and 3) provide a learning process that allows (1) to become (2) (Fodor, 1980; Carey, 2009). The section will proceed in this order, starting with the “adult state,” or what we may consider the goal of learning, and turning to linguistic data that may bear on the initial state, and ending on non-linguistic experimental findings inspired by paradigms in comparative psychology. This will serve to set up a discussion of (3) and difficulties in developing theories of such learning processes in the next section. (1), of course, is approached by the experimenter such as to reduce the likelihood of false negatives (e.g., failures due to performance) and false positives (e.g., successes due to a confound resulting in alternate strategies) so though designs and findings will be summarized, the present argument will not turn on whether the experiments were well-designed, but instead primarily on the interpretation of the findings.

Adult State

Negation in logic

In the final state, negation allows us to easily talk about error, absence, and other lacunae. Without it, we would have to be omnipotent and omniscient to hold a decent conversation (Bacon, 1620). While negation is formally definable, this is not to say that there is an uncontroversial definition. As logic is an active area of research, there are various approaches that make different assumptions about how negation works (Harman 1986a; for a review, see Gabbay & Wansing, 1999 and Wansing, 1996) for a host of attempts to formalize it within intuitionism. In first-order logic, though, which highly influenced contemporary semantics, it can be represented with a simple truth table (Table 4).

P	¬P
0	1
1	0

Table 4: The truth table for negation.

That is, contradictory negation is an operation that inverses the truth value of a proposition, such that if P is a true statement, ¬P contradicts P, and is thus necessarily false. If, on the other hand, P were false, then ¬P would be true (see Categories in Aristotle, 1975; Horn, 2001). Aristotle considered negations to be a form of opposition, a broader class of relations which includes what we now typically consider non-negative relations, like correlations (e.g., double and half) and privatives (e.g., bald when hair is expected). He referred to logical negation as "contradiction" and used it to define the logic of opposition:

> The Law of Noncontradiction: (LNC) A proposition and its negation cannot be simultaneously true.
>
> Law of the Excluded Middle: (LEM) For any proposition, either it or its negation is true.

A statement and its contradictory cannot be true or false at the same time. However, between a statement and its contrary, it is only required that they are not true at the same time. For example, the contradictory of "Crows are black" would be "Crows are not black," and must be true in any case where "Crows are black" is false. However, a contrary to "Crows are black" might be "Crows are white," and while both statements cannot be true at the same time, they could both be false (e.g., if crows were pink). Horn (2001) suggests that another way of conceiving the difference between contradictories and contraries is that while both are governed by LNC, only contradictions are governed by LEM. Contradictory negation (truth-functional negation, logical negation) is what is typically referred to when negation is discussed because it is ubiquitous in classical semantic theories.

Negation in Natural Languages

In English, standard negation (i.e., phrasal or sentential negation, of the sort “That is not a dog.”) must be commanded by auxiliaries. This means that to transform an affirmative statement (e.g., “She opened a can of worms.”) into its negative counterpart, an auxiliary must be added, which additionally takes the tense (e.g., “She did not open a can of worms.”; Horn, 1989). Standard negation is also often expressed in contracted forms of not (e.g., *don’t*, *won’t*, *shan’t*), but beyond standard negation there is also no. No can serve as a quantifier or stand-alone anaphorically in response to a question. Typically, not is treated as the natural language analogue to logical negation, with no serving a range of communicative functions (i.e., speech acts) that are more specific, from communicating displeasure, Rejecting an offer (“No veggies.”), quantifying and remarking on Nonexistence (“No water.”), Prohibiting (“No running.”). In addition to being more contextually dependent, these functions seem to lack the inverting character of truth functional negation (e.g., two prohibitions do not cancel out), and so many assume it is not necessary to represent such thoughts. The only time no is uncontroversially considered to involve logical negation is when it’s used anaphorically to Deny a polar question (A: “Is that yours?” B: ”No.”). The speech act of Denial is often equated with truth-functional negation because it seems to share its inverting qualities (e.g., denying a denial cancels out) and can also be used in most domains (e.g., not only those pertaining to desire or permission). However, it’s unlikely that every adult utterance that can be given a truth-functional analysis is reducible to a Denial as there are many cases where negation is used in reasoning that do not amount to a Denial, even if negation may is commonly used to to do so (Wason, 1965; 1972; Johnson-Laird & Tridgell, 1972). For example, someone who utters “If Paris is not the capital of France, my itinerary is in trouble,” is likely not in a position to deny the truth of “Paris is the capital of France” with certainty, despite using it in a formal context (an “if… then” conditional; Frege, 1919). Despite this qualification, Denial takes on special relevance developmentally, as early negative utterances are typically simple and can thereby serve as a bellwether for formal negation rather than as entirely identical with it.

Initial State

Linguistic findings

Studies of child production have found that though children use no fairly early, with most parents reporting that their 16-month-olds children regularly produce no when surveyed (Frank et al., 2016), they do so only in some particular ways. Typically, they use no early on – to Reject an offer or command, or, more rarely, to talk about Nonexistence. Children do not use no to Deny propositions until about 2.5 years, around the same time when they also begin using not more regularly (Bloom, 1970; Choi, 1988; Nordmeyer & Frank, 2017). These production patterns have been interpreted in both ways that assume a broader conceptual discontinuity as well as not. As an example of the former, some have proposed it is the result of constructing logical negation from a simpler more general kind of negation (e.g., meta-linguistic as in Hummer et al., 1993; Bloom, 1970; or outward rejection, where a mental state is at odds with a state of affairs, as in Pea, 1980). As an example of the latter, it has been suggested that children may be limited in their production for syntactic reasons (e.g., being limited to two-word stage as suggested by Bloom, 1980 or not yet learning the requisite affirmation-to-negation transformations (McNeill & McNeill, 1967) However, later studies have found that children also do not comprehend Denial until about the same age as they produce it, which cannot be explained by these performance factors related primarily to production. These studies involved search tasks where a target (a ball) was covertly placed into one of two places (bucket/truck). The experimenter would check one of the places and tell the child that the object was not there (Austin et al., 2014; Feiman et al., 2017). It’s not until around the same time that children begin producing not (about 2.5 years) that they also succeed at this comprehension task. However, French-speaking children as young as 18-months-old have been found to understand negative sentences on a simple habituation task (Carvalho, Crimon, Barrault, Trueswell, & Christophe, 2021). In this study, children were presented with a word learning task designed to reduce any potential difficulties children might have with negation due to their limited inhibitory control abilities (i.e., to control for potential performance factors masking competence). Children were tasked with learning a nonsense word (e.g., bamoule) and were exposed to it in either a noun frame (“Oh, look! It’s a bamoule.”) or a verb frame (“Oh, look! It bamoules.”) while watching a penguin complete one action (either spinning or cartwheeling). When exposed to it in a noun frame, participants should learn it roughly means “penguin,” whereas when exposed to the verb frame, they should learn it roughly means whichever action they were exposed to. After habituation, they were then exposed to a trial in which they heard a negative statement (e.g., “Oh look! It’s not bamouling.” vs. “Oh look, it’s not a bamoule.”) while watching the penguin complete a different action than at training. Assuming that children in the verb-frame condition learned that bamoule means to perform one action (e.g., cartwheeling), then hearing that “it is not bamouling” when completing another action (e.g., spinning) should remain true in context (cartwheeling is not spinning) and thus children should not display signs of confusion (i.e., increased looking times). However, assuming participants in the noun-frame condition learn that bamoule means penguin, then as they are still shown a penguin when hearing “It is not a bamoule,” they should exhibit increased looking times during negative trials. This is in fact what they found, along with additional evidence that 24-month-olds are able to learn word meanings from negative contexts. Thus, when Executive Function demands are minimized, recent evidence suggests that children exhibit understanding of negative sentences earlier than believed - significantly closer to attested productions of no. However, before children learn the meaning of no there must be some period (however brief) where they have the concept (Gomes et al., 2023) without the label, and thus, as Feiman et al., (2017) mentions in the discussion, comprehension studies leave open whether any potential failure of earlier participants are simply because they cannot represent logical negation, or because they just have not learned that certain words in their language map onto it. This limitation is inherent to linguistic tests of negation, and so non-linguistic experiments are necessary to make a stronger claim at the earliest ages.

Non-linguistic findings

Non-linguistic studies into negation have taken one of two approaches: attempting to teach children to generalize to a rule that includes negation, and gauging their ability to use an inferential rule, particularly the disjunctive syllogism (A or B; not A; therefore, B; see Cesana-Arlotti et al., 2018 and Feiman et al., 2023). Hochmann and Toro (2021) attempted to teach 11-month-olds to differentiate between syllable strings where all syllables were the same (AAA) and strings where all were the same but the last (AA¬A). However, infants only showed success when the strings were certain lengths (AAAA and AAA¬A), and while the authors suggest this may be due to cognitive load, others have argued this may be evidence of a non-general function (and hence, not full-fledged negation). On the other hand, methodologies testing infant's ability to use the Disjunctive Syllogism have met with some success in 12- and 19-month-old infants (Cesana-Arlotti et al 2018; hereafter referred to as the "two-cup task"). Participants were shown two items and a cup. The items were designed so that they crossed semantic categories, but also had similar top halves visually (e.g., a yellow dinosaur with purple spikes and a purple flower). An occluder comes up and obscures the items completely, after which the cup goes behind the occluder and scoops up one of the two items (which look the same in the cup). The occluder is lowered, and the non-scooped item is revealed. Finally, the cup is lowered to show either the item that was not behind the occluder (consistent condition) or the item that was behind the occluder (inconsistent condition). This found that participants in both age groups looked significantly longer during inconsistent outcomes.[15] Prior to this,

[15] As well as a follow-up controlling for the possibility that infants were just responding to seeing the same item twice in the inconsistent condition.

success in a Disjunctive Syllogism-based class had only been found at 2.5 years (Mody & Carey, 2016). However, that study employed an alternative method where 1) children were asked to make a decision (pick a cup the sticker is in) and were given more alternatives (4 options rather than 2). Criticism of Cesana-Arlotti et al (2018) primarily centered around the former, in that constraining the option space to two alternatives allowed for nonlogical ways of completing the task, and hence the data does not constitute uncontroversial evidence of logical negation.

Jasbi et al (2019) argued that the two-cup task could be solved with a mechanism that children are known to employ early on: object tracking (Wynn, 1992; Xu & Carey, 1996), and this may be why they succeed at this version of the task but fail at versions with more objects (Mody & Carey, 2016). Infants may be creating a mental model of the scene, where locations and objects are tracked. Upon seeing two items, an entry for each (with parameters: item name, location), and when the occluder goes up, their "location" parameter changes from null to "behind occluder." When the cup scoops something, both item's location parameters are deleted, and a new location entry "inside cup" is added, but the item name is empty since the visible portion is ambiguous. When the unscooped item is revealed, they can switch its location to "behind occluder." That results in the other item lacking a location, and a location ("inside cup") lacking an item, which could then be "bound" to each other. Feiman et al. (2022) conducted additional experiments with this alternative strategy in mind and found evidence that 17-month-olds were able to perform above chance on tasks in which an object-tracking strategy could not be used to complete the task successfully. Thus, we can say that though there is some very early evidence at 11- (Hochmann & Toro, 2021) and 12-months-old (Cesana-Arlotti et al., 2018; for additional work suggesting success at 14-months from the same team, see Cesana-Arlotti et al., 2022), the earliest evidence children can uncontroversially negate is starting at 17 months (Feiman et al., 2022).

Learning Proposals

From Ethology: Domain-Specific or Partial Function

Regardless of when children can successively employ Disjunctive Syllogisms, a related question is how they gain the ability to do so. That is, after all, the rub of concept learning. Cesana-Arlotti et al (2018) interpret their evidence as suggesting “primitive logical abilities,” “precursors of logical reasoning,” and that “the development of reasoning abilities builds on a natural logical foundation.” The authors seem to imply that some aspect of logic is innate, but they are noncommittal in regard to how/what parts of the Disjunctive Syllogism are. At one extreme, either Disjunction and/or Negation (and machinery required to combine them with propositions and arguments) are innate, however, a less extreme interpretation would be that however the task was solved may still involve representing negation (i.e., perhaps object tracking requires negating).

Contrariety has been referred to as “protonegation” in the animal cognition literature, which is interested in the origin of concepts evolutionarily rather than developmentally (Bermúdez, 2003; Bohn, Call, & Völter, 2020; Völter & Call, 2017). However, their interpretation of contrariety seems inspired by the concept of immediate contraries (akin to “opposites,” as suggested by Aristotle (Aristotle, 1975) which do in fact end up obeying the LEM because together they categorize the space of possibilities (e.g., even if we do not yet know which is which, a number must be EVEN or ODD, and a person must be DEAD or ALIVE). As such, they distinguish contrariety from logical negation by domain-generality (e.g., being reusable through cognition with a symbol). Bermúdez (2003) argued that without logical negation, animals could learn that pairs of predicates are (immediate) contraries (PRESENT - ABSENT) and that this could allow them to seem to exhibit the use of logical negation in contexts where they have learned two properties are contraries. However, it’s unclear, as McDermott-Hinman & Feiman (2022) point out, how one could learn that two things are contraries without being able to represent contrariety in the first place, and Bermùdez's treatment does not meaningfully resolve this problem.

1. P(E1) → ¬P(E2)
    1. "If the gazelle is present at the watering hole, the lion is not present at the watering hole."
2. P(E1) → contrary(P(x))(E2)
    1. "If the gazelle is present at the watering hole, the lion is absent at the watering hole.
3. P1(E1) → P2(E2)
    1. "If the gazelle is present at the watering hole, the lion is absent at the watering hole.
    2. "If the sky is smoky, the fire is lit."
    3. "If there is a bear, it shits in the woods."

Bermúdez (2003) argued that an observer that could be described as learning a conditional with a negative consequent like (1), could also be described as conditionally relating a predicate term and its contrary as in (2). Contrary pairs could be learned by representing a conditional as in (3), and considering a situation where the individual in the antecedent is predicated by the (P1(E2): "The lion is present at the watering hole.") to determine whether it is a protonegation. While one could describe the situation in question this way, it is not a valid argument, as many causal relationships (3b) and correlations (3c) are describable in the same form but are certainly not contraries. Mediate contraries are rare in being totally negatively correlated, so if such a mechanism is to work, it needs to distinguish between rare and impossible co-occurrences or mediate contraries would need to be innate. Even so, Bermudez (2002) goes on to suggest that contrariety may not figure in the description of language nor logic (and hence is not present in the final conceptual state, though it may continue to figure in nonlinguistic processes), but that it could serve as the foundation for logical negation when it is transformed into a domain-general operation. Bohn et al. (2020) argued the transition out of protonegation evolutionarily may require the development of propositional thought, which could be the result of social pressures (i.e., needing to communicate and coordinate with other agents).

Developing Contrariety

Feiman et al. (2020) however, argued that this view of the relationship between contradictories and contraries conflates two ways in which a representation can be implicit: domain-specific and partial function. If a representation is domain-specific, it is encapsulated within particular computations or contexts (e.g., Rejection (NOT-WANT) and Non-existence (NOT-EXIST) uses of negation) and not available for use outside them. For example, perhaps the object-tracking mechanism can manipulate mental models with an operation that employs logical negation, but such a use of logical negation may be inaccessible outside of this mechanism. If a representation is a partial function, it figures in the definition of some target concept (e.g., negation) but is not sufficient for possessing it.

McDermott-Hinman & Feiman (2022) fleshed out a proposal for how a partial function account could apply to negation, using an account of contrariety rooted in the notion of mediate contraries (e.g., HOT and COLD). They argue that children are capable of contrariety as an operation on statements which obeys the LNC, and that acquiring logical negation requires constructing a concept that uses contrariety to represent LNC and further constrains it so it is subject to LEM. They note that while a statement ("You are in New Jersey.") may have only one contradiction ("You are not in New Jersey."), it can have many different contraries, or statements with which it can be simultaneously true ("You are in outer space," "You are in France," and so on).[16] As such, they suggest the negation of p could be defined as the set of all the contraries of p ($\sim p := \{q \mid contrary(q, p)\}$). They therefore do not consider contrariety protological, but perhaps do consider it to be aptly described as protonegation. While logical negation would not be innate under this approach, it therefore suggests a weak discontinuity[17] as other aspects of logical

[16] Just as a term may have many mediate contraries, though BLACK and WHITE are often seen as opposites, any other color is also a contrary of WHITE (e.g., nothing can be RED and WHITE at the same time, but can be neither (e.g., when in between or any other color).

[17] Other theories, which posit soft logical discontinuity, tend to tie the conceptual limitation to an ability that children develop (e.g., inhibition, Theory of Mind). These approaches need not commit to representational discontinuity of negation unless one identifies the ability with the concept (e.g., inhibiting IS negating), or if one does not ascribe to the performance/competence distinction.

construction (definition, set theory, contrariety) would need to at the very least be in place very early on and thus the terms in the adult stage are definable in the initial state. However, even a weak individuation discontinuity account runs into the issues introduced in Section 2, so we will now turn to the issues with this account.

Not learning NOT

The protonegation definition provided by McDermott-Hinman and Feiman (2023) suggests that the delay in word learning is due not just to e.g., access discontinuity but also because the concept in question is complex (~p := {q | contrary(q, p)}) despite corresponding to a monomorpheme (not). Part of the appeal of definitions, under the classical theory, is that definitions provide insight into the learning process, but it's unclear how to translate this into a learning story. On the one hand, this is because as suggested, there are open questions about contrariety itself, and on the other, because it risks reintroducing holism. I will briefly discuss some issues with contrariety in logic before briefly relating it to the issues discussed in Section 2.

Under the classical theory, primitive concepts are more immediate than complex concepts, and complex concepts can be more or less complex (e.g., . That is why it would predict that children learn MAN before BACHELOR, MAN is easier to think of spontaneously (if you are a child) than BACHELOR. If CONTRARY is more primitive than NOT as this definition might suggest: what would it predict about word learning? By the present account, contrariety (CONTRARY) would be applicable and roughly useful most of the time by the time truth-functional negation shows up. Instead of understanding someone's utterance about it not being cold as NOT(COLD), one could understand it as being about CONTRARY(COLD). Yet, in English, at least, though there are monomorphemes (e.g., *un-*) that some argue correspond to contrariety, they are bound monomorphemes, and are additionally rarely attested in early word learning (see Section 5.1 in Horn, 2001 for a summary of such views). As such, it appears there is no early-arriving word which corresponds to domain-general contrariety and which remains in use primarily as such, despite the fact that proposing it as a primitive would in most cases suggest early acquisition (I.e., since only mapping is necessary).

These considerations might suggest that contrariety, even when supplemented with set construction, is not sufficient for capturing uses of negation. It's important to remember that there are different notions of contrariety within term logic,[18] but even more importantly, there is doubt as to whether contrariety is formalizable within propositional logic to begin with (Englebretsen, 1974; Geach, 1975; Horn, 2001). This would pose a problem for the present definition as, on the one hand propositional logics cannot represent contrariety, and on the other, term logics cannot represent quantification. Second, non-classical logics provide interesting insights here, as such early logics did not build in the Law of the Excluded Middle (for examples, see Brouwer, 1977; Heyting, 1980; Horn, 2001), and as such the question of how it could be constructed, approximated, or otherwise explained away while maintaining other commitments about intuitionism (i.e., anti-realist and constructible) has been an active topic of discussion (Gabbay & Hunter, 1999).

Alongside contrariety (Tennant, 1999), logicians have tried to construct the LEM out of: absurdity (e.g., 1=0: Dummett, 1993; Valtonen, 2017), inconsistency (e.g., something is a negation of p if together with p, unwanted formulas can be derived; Harman, 1986ab; Gabbay, 1987) and falsity or disproof (e.g., something is a negation of p if it actively supports the falsity of p; Nelson, 1949; Wansing, 1991; Wansing, 2001). As these disagreements suggest, there remains no consensus on the definition of negation amongst non-classical logicians either, and so the suggestion that negation is in principle definable should be amended to clarify that this does not therefore mean we have reached any consensus. Although Wansing (1999) convincingly argues that every case of negation as inconsistency and contrariety can be proved to be cases of negation as falsity, while the converse does not hold, suggesting it covers some

[18] For example, ALIVE and DEAD are immediate contraries, something that can be either cannot be both or neither, while COLD and HOT can be mediated contraries as something that can be either can neither but not both.

aspects of negation ("strong negation") they lack.[19] Unfortunately for present purposes, it is difficult to give an inferential account of falsity (Hand, 1999), as it is affected by psychological notions (like immediacy in Harman (1986a)) or social-pragmatic notions like obligation (e.g., do not say false things). And, like contrariety, none of these other formalizations (absurdity, inconsistency, falsify) suggest a precursor which is lexicalized early on.

This is not to say that negation as contrariety is certainly incorrect, but that as a formalization of truth-functional negation it is currently lacking and the proper inferential role treatment of negation remains unsettled. As a consequence, there are in fact a staggering number of ways to conceive of negation as a complex concept and thus attempting to account for how negation is individuated from how it is accessed will risk holism. These non-classical proposals produce identical behaviors in a range of contexts, but diverge in others, so how do word learners converge on not just the correct inferential rules, but the correct notion of negation to begin with based on a limited amount of ambiguous experiences? Furthermore, it's uncertain whether it's possible to constrain the notion that only monomorphemic logical concepts can be complex, which risks introducing holism for non-logical words which may similarly be given inferential accounts (e.g., prepositions like ABOVE and UNDER). What this and other learning stories of this sort are missing is not just the right story about how one constructs one concept out of another, but about what would compel one to construct it in some particular way if one were not already so inclined. This remains as difficult for logical words as non-logical words, though this may be harder to recognize given their interdefinability.

Characterizing the input of non-linguistic experiments is complicated, as the input in this case would not be linguistic but perceptual, which is less constrained. As the adage goes, a picture is worth a thousand words and that's exactly the problem (Gleitman & Gleitman, 1992). This issue is less pressing for words that have fairly straightforward extensions, like DOG, but

[19] To be clear, this is not meant to suggest that this argument is final, as negation as falsity cannot accommodate contraposition (A->B; not(A)->not(B)) though some consider this critical for a complete understanding negation (Restall, 1998; Lenzen, 1996).

NOT and other logical concepts do not pick out categories of things in the world. In these cases, one cannot rely solely on characterizing visual aspects of the scene as one might for DOG (e.g., by similarity of pixels or features). Instead, experimenters must develop tasks that rule out alternative strategies, and it's uncertain how these tasks should be characterized in terms of inputs, and therefore which strategies include negation and which do not. There is unlikely to be a characterization of a task or scene that can be guaranteed to be apt and thus likely to be shared by individuals, even adults. This, in a sense, is the problem of access, which is why one particular representation of a task would be more likely to be accessed over others given an underdetermined input (especially when one does not know much else about the world).

Given the potentially non-formal nature of contrariety (i.e., from knowing that P is false, one cannot know that contrary(P) is false, nor what predicate it would use), and the low likelihood of testing potential processing predictions, this learning mechanism is unlikely to be adequate for describing how one overcomes early representational limitations. Additionally, it gives little suggestion as to why we see the patterns we do - why do kids use logical negation around 2 years rather than 3 - as well as to what sorts of learning experiences would activate or assist this process. As such, it, at the very least, requires further refinements. Ideally this would involve establishing predictions it would or would not make about processing or learning instances, which may indicate why kids tend to learn logical negation around 2 years. However, even if such refinements are made, it's important to acknowledge that these claims arose out of particular interpretations of behavior in non-linguistic tasks. I will argue that these are theoretically loaded, and as such, one may not need to provide a learning mechanism for concepts at all.

Summary

While it was believed that not was learned much later than no, recent work suggests this may not be the case, as French-speaking children as early as 18 months old can understand negative sentences when the task design reduces Executive Function load (de Carvalho et al., 2021). Attempts to probe nonlinguistic ability through teaching children “grammatical rules” about an artificial language which require negation (e.g., AA¬A) has met some success as early as 11 months (Hochmann, 2020; Hochmann & Toro, 2021), before children are regularly using no, but at this age their performance depends on contextual factors (e.g., number of syllables), which should be irrelevant to logical negation (Feiman et al., 2023). While there has also been some success around this age in tasks which test infants ability to reason using the Disjunctive Syllogism, namely the two-cup task (Cesana-Arlotti, 2018), others have noted that this design did not rule out an alternative strategy (object tracking) which would not require negating anything (Jasbi et al., 2019; Feiman et al., 2023). When rendering this alternative strategy ineffective, success is only found around 17 months (Feiman et al., 2022; Gomes et al., 2023), which would importantly still place it earlier than previous estimates (Feiman et al., 2017; Austin et al., 2014), and just a few months after children are reported to say no. The central question of this paper is not when exactly children learn that not means NOT, but whether doing so requires doing more than mapping not to NOT. Thus far, it seems like not may not be learned so far after no, suggesting the same word learning mechanism may explain both (and thus that one does not require recourse to concept learning). This is perhaps good news, as even the most developed concept learning theories for truth-functional negation run into issues with holism. This is because it runs into issues as a purely formal description, and therefore attempting to translate it into a learning story would vastly complicate the learning process. Given this ambiguity about how to represent negation and the early acquisition of not alongside other early words, the argument for Mad-Dog nativism remains relevant for logic words. The next section will attempt to sketch out how Access Discontinuity in light of this.

Access Matters

> *He once greeted me with the question: 'Why do people say that it was natural to think that the sun went round the earth rather than that the earth turned on its axis?' I replied : 'I suppose, because it*

looked as if the sun went round the earth. "Well,' he asked, 'what would it have looked like if it had looked 'as if the earth turned on its axis?'

- G.E.M. Anscombe, 1967

You can think of a cow (and in doing so token COW) for a variety of reasons: from seeing one, hearing a "moo," or, as in this sentence, from hearing/reading the word cow. Some modes of deployment may seem core to a concept, but there's no reason to take any stimulus as equivalent to the (conceptual) response it evokes. Most would uncontroversially agree that even though they may not be thinking a COW thought at a given moment does not mean that they are incapable of it. As Beck (2014) noted, just because a concept is manifest does not mean it is currently "active" or deployed by the conceptual system. Clearly by manifestation, as discussed in Section 2, Carey (2014) does not mean that adults are at all times actively conceiving of, e.g., number or density, but only that the potential for such representation exists. It seems fairly uncontroversial to claim that we seem to have manifest concepts that we deploy as required (or as we please) but that are otherwise not actively represented. This may seem like a small point, but it serves to highlight that having a concept and when one uses it are two different things. This raises the question of how we, as scientists, could tell that a concept is missing rather than inert under some presentation (e.g., an experimental condition). In other words, a manifest concept need not be deployed, and there are plenty of failures to deploy a concept that we do not equate with non-possession (e.g., certain laundry detergents remind me of playing Super Mario 64, if they do not for you, fret not! You can still think Mario thoughts). The recourse to manifest concepts as the only proof of concept possession is not theory neutral, and instead belies the assumption that concepts should be similarly accessed in addition to being identically represented (i.e., by nomic connections) and that mastering a concept requires this dual achievement. This assumption, is, in part due to the same mistake made by the British empiricists - to seek an explanation for the structure of reason outside of the mind. Adult speakers of a language may think DUTCH when exposed to a painting, but they may just as easily think BEAUTIFUL, HIDEOUS, and CLASHES WITH THE WALLPAPER (Chomsky, 2013). There's no reason to assume routes of semantic access are or should be identical between two individuals, especially if they differ greatly in age. First I will briefly describe Fodor's notion of semantic access before outlining how a discontinuity in semantic access (not individuation) may account for the range of findings that motivate individuation discontinuity theorists.

As noted in Section 2, Fodor's attempt to steer clear of meaning holism is part of what led to his endorsement of information semantics, but it alone cannot explain why some concepts are difficult to access, despite being individuated by the learner. If concepts are merely nomic connections between their tokenings and the things they token, then why do children learn NO before NOT? The answer is fairly straightforward under a model where NOT is not atomic, and thus more complex than NO, but if NOT is similarly given as I have argued Mad-Dog Nativism still suggests, then why do we not map it to a word as quickly as NO? It additionally does not explain "how… the instantiation of doghood causes tokenings of DOG." Thankfully, Fodor's answer to the latter issue may provide a hint to the former. He argues that there are modes of semantic access for concepts. In short, there are things, like seeing a dog in the middle distance, that allows a learner to semantically access the concept in question. The list of ways a concept may be accessed is open-ended, since one could e.g., condition a junior faculty member to think TENURE every time they see a carrot.[20] Understood this way, a word is also a route of semantic access for a concept. Under this account, to learn a concept is to become "nomologically locked" to the properties it expresses, or proposing a concept for the meaning of a mono-morphemic word that one can then semantically access through various routes. It does

[20] A commenter expressed doubt that it is truly open-ended. For example, would it be possible to condition this assistant professor to think TENURE every time they breathe in? Or, sub in some other activity typically considered to be involuntary (and perhaps that is always occurring, every time they see light. There are likely constraints, and so to say it is open-ended is not meant to imply that anything at all could serve to access a concept, even with conditioning, But that one could establish chains of conditioning, theoretically allowing the list to be infinite (while simultaneously not allowing for every single action to be a potential route of semantic access).

not, however, guarantee the concept will be regularly accessed under every condition associated with its access conditions, as each route of semantic access may be subject to different performance conditions (e.g., language, audibility, visibility, primacy, recency, cognitive load).

Perception may head the list of semantic access for various concepts, but importantly, the concept one entertains as a result is not as a consequence defined by the mode it was accessed by (e.g., it is not a list of perceptual features; Halberda, 2019). Undoubtedly, the list of routes to semantic access changes based on experience, that must be part of what word learning is in itself – that you may now access a concept based not on tokenings of e.g., members of it in the world, but by sound waves. In such a case, words are an example of semantic access by convention (i.e., arbitrariness of the sign, nothing in the sound “dog” or “cachorro” that triggers DOG). It may be straightforward to spot conventionalized routes of semantic access from non-conventionalized routes in some cases, but suddenly, Quine's critique of analyticity raises its head. Determining whether a route is non-conventionalized requires many auxiliary assumptions about not just what information is in the environment, but what a child’s initial conceptual state is. Thankfully, that is the task of the scientist and not of the child, despite the analogies drawn between the two. Children do not need to know whether something is conventional or not or why, they merely learn using the routes afforded to them, until they acquire the systems that allow them to more easily access other concepts. The gist of Access Discontinuity is that children are not scientists, but overworked opportunists with a limited budget on a quick turnaround.

New sources of information affect conceptual access

If access discontinuity is correct, it would strongly predict that adults would behave like children along a range of word learning tasks under limited information conditions,[21] and in fact

---

[21] Assuming the information in question helps in accessing the concept.

adults do. The Human Simulation Paradigm was developed to test precisely this question, but in response to the claim that nouns predominate the early vocabulary because they are conceptually simpler as they are syntactically simpler (Gillette et al., 1999). In this approach, adults are placed in the shoes of children at various stages in their linguistic development by giving them access to the different parts of information present in natural parent-child interactions. In one condition, designed to mimic early learning by visual observation only, participants watch a short video of a natural parent-child interaction which is silent other than a single beep and are tasked with guessing what the parent uttered at the beep (either a noun or a verb). Participants were much more likely to correctly guess which nouns were uttered when given no additional information, but in other conditions, as when they were given syntactic information (even when it was nonsense, like "The dax GLORPS that zeb is flerpy."), participants were much more likely to correctly guess which verb was uttered (Gillete et al, 1999; Snedeker & Gletiman, 2004) despite already having been told whether the upcoming target word was a noun or a verb in all conditions. This effect was especially drastic for think (for a review, see Gleitman et al, 2005). Gillette et al, responded to conceptual change accounts by proposing that instead of a change in conceptual ability, there is a change in the kinds of information children are able to make use of. Though kids may always receive linguistic input with the syntactic information they need, they cannot yet parse them into structures at the earliest stages of word learning. Thus, they proposed information change as an alternative to conceptual change as the principal driver for the patterns in question.

This approach was also extended to address conceptual discontinuity claims concerning negation. It's been frequently shown that children produce no earlier than not, and that the sorts of things they use them to do seem simpler than the true logical meaning of negation. This has led many to suggest that this is because children cannot represent the logical function of negation proper and so the early uses do not map onto it, but onto simpler concepts like rejection and nonexistence (Bloom, 1968; Choi, 1988; Nordmeyer & Frank, 2014; Feiman et al, 2017). Some have even explicitly suggested that truth-functional negation may be constructed through Quinean Bootstrapping, a hypothesized learning mechanism not rooted in hypothesis testing suggested to explain word learning patterns for numbers, using these simpler concepts (Feiman et al, 2017). However, recent research has similarly found that, when shown a muted video of naturalistic uses of negation by parents to their child, adults struggle to infer that negation is used from denials (truth-functional negation) compared to the simpler kinds of negation, unless provided with additional information (e.g., everything but a blank for the negator, "That is ___ yours"; Gomes et al, 2021).

The Human Simulation experiments highlight that a difficulty remains for language learning even when you have all the concepts: the mapping problem is not just about mapping a concept to a label, you have to infer that someone is using a particular meaning for an unknown label or description, and this may be harder for some kinds of words than others. In other words, it's not just about being able to think about e.g., ONE, but doing so under the same conditions as others that speak your language (using similar routes of semantic access). The information change proposal argues that through learning language, children are able to use facts about the language to assist them in learning word meanings. Syntactic bootstrapping is a key mechanism developed by Gleitman and colleagues (see Gleitman et al, 2005 for review) and has been supported by non-HSP experimental work (Naigles, 1990; Yuan & Fisher, 2009). While they have not explicitly drawn the connection between syntax and semantic access, they refer to syntax as a "cue" to meaning. Given the definition of semantic access and the examples above, there's no principled reason to think that this difference in terminology belies theoretical incompatibility. The most important aspect of Fodor's proposal of semantic access is that no route is itself constitutive of the concept that is accessed, and this is satisfied by "cues" just as well. Gleitman and colleagues have not claimed that the syntax of e.g., think simply is the meaning, merely that it helps learners think someone else is thinking THINK (or credal verbs more broadly). As such, information change may be fairly construed as a combination of conceptual continuity (no conceptual growth) and access discontinuity (routes of semantic access change through development). However, access discontinuity extends to a greater range of patterns in word learning than just those for which Information Change has been argued (mostly verbs, credal verbs, and negation thus far). Still, Access Discontinuity alone does not guarantee an answer to our question, though it may be motivated by the same set of experiments Information Change was advocated for independently. For it to generate the patterns we see additionally requires a better understanding of the task of word learning as well as the potential mechanisms involved.

Knowing is only half the battle

Access in linguistic contexts

Indeed, as mentioned earlier, perception may head the list of semantic access, but this fact cannot guarantee that a word will be learned since word learning does not consist solely of explicit and unambiguous labeling events. Instead, children must navigate a referentially and semantically ambiguous world and learn across situations to learn words. One potential class of early word learning mechanisms argues that children keep track of many meanings at a given exposure, and are able to compare this with the meanings they entertained for that same label under other exposures (global cross-situational learning models). They can compare the distribution of hypotheses (cats are around a lot when cat is uttered, but not dog) to find what was the most likely referent, and thus determine the meaning (Yu & Smith, 2007; Smith & Yu, 2008). Another class of proposals argues that children only track a single hypothesis from an experience, the one they felt strongest about, and test it on following exposures (local cross-situational learning models; this global/local distinction originated in Stevens et al, 2017). If they are able to confirm that hypothesis, they strengthen it and continue testing it. If they do not, they propose a new hypothesis based on the current exposure and either forget the old one (Trueswell et al, 2013) or penalize it (Stevens et al, 2017). Under these models, there is no comparison of hypotheses (e.g., determining their similarity or structural proximity) to allow for the generation of a new hypothesis: instead, the best guess at that moment is hypothesized and it either is the same as the previous hypothesis or not. An upside of local models is that they are less computationally demanding, they do not require every potential hypotheses to be stored, and they do not allow for changes in one hypothesis to affect changes throughout a network of hypotheses as global models do, while still allowing for the resolution of referential uncertainty in early word learning. Thus, it's easier to conceive of a memory-limited learner being able to "make the most" out of this word learning approach, since it drastically reduces processing cost and what is stored in the memory.[22]

The task of early word learning, as children approach it, requires making quick guesses in real-time. Due to their more limited cognitive abilities (e.g., memory or inhibitory control) and world knowledge, certain hypotheses would just not tend to be formulated early on as they would require more information to trigger within the timescale required for consideration in word learning (e.g., wholes before parts (perhaps because you would need to start thinking about which parts), basic level before others (perhaps because you need motivation to propose a hierarchy)). Children propose what comes quickest, e.g., some characterization of the object when hearing a label in isolation, and do not compare hypotheses. As they acquire knowledge of new systems, they are better able to access a wider range of concepts. As noted in the previous section, language is one such system, and its acquisition greatly improves the learner's ease in accessing concepts and manifesting their total expressive capacity. Without it, accessing certain concepts, like THINK and NOT may prove difficult, but not impossible. While language may aid the mapping problem, coordination of meanings between speaker and listener is no simple feat. Even for the earliest words there seem to be contexts that are very strong learning instances, and this relates to the time signature of ostensive reference and joint attention (Baldwin, 1991; Gleitman & Trueswell, 2020; Trueswell et al, 2016). For example,

[22] One might note that the desire to avoid global models for word learning stems from a similar place as the desire to avoid meaning holism, but this is left as an exercise to the reader.

experimentally decoupling a labeling event from drawing attention to an object (i.e., picking it up and moving it) decreased performance at test for children younger than three years years of age even when the displacement varied by ±6s (Trueswell et al, in prep), and altering highly informative HSP vignettes by ±1-2s had a similar effect for adults (Trueswell et al, 2016). Since even for “easy” words, joint attention and synchrony is important, it’s not unreasonable to assume the same holds for all words.

Under this view, even if children are as conceptually equipped as adults (e.g., they “possess” but do not “manifest” every concept an adult can), they will exhibit patterns in word-learning (and learning more generally, as much as it interfaces with word-learning) that may look like the result of a more rational, theory-building process, but is instead a consequence of tracking simple, unstructured hypotheses across exposures. As such, any behavioral discontinuities may be due to children simply not immediately accessing negation in such tasks, and relatedly may not do so within the time window of early word learning. However, by learning the syntax of their language, negation can be more immediately accessed in those contexts where negators are uttered. By learning the label, one can thereby learn other contexts where one’s language community finds it useful to negate, and thereby access it more regularly even in the absence of negators being uttered.

Access in non-linguistic contexts

While Access Discontinuity accounts of word learning, like Information Change, have primarily focused on investigating when new sources of evidence become accessible to learners, they additionally gesture towards a broader, non-performance-based account relevant to the interpretation of non-linguistic behavior. Children are not just gaining access to new sources of information, but also learning more about the world in light of the information at hand. As the last section suggested, by learning a word, one is better equipped to learn about the contexts in which people typically find it relevant. This is especially relevant for logical words, as it's unclear whether their wide-spread utility is so immediately obvious to children. That is, thinking about something does not solely require the ability to use it, but think it's relevant to accomplishing some ends. A failure to find evidence of a concept, logical or otherwise, in a particular context does not mean much without an argument as to why that particular context should be meaning constitutive over others. The basic thrust of this paper is that what advocates of competence discontinuity (Feiman et al., 2023) need, but do not yet have, is a story as to why the tasks that they claim can be solved with alternative strategies (e.g., the two-cup task in Cesana-Arlotti et al., 2018) do not themselves employ or assist in employing negation (e.g., but instead use solely object tracking). The present approach of controlling for alternative strategies for claimed successes risks ruling out the very cases most relevant to early word learners, and therefore be of limited use to word learning research. This is because the contexts in which negation may be more readily accessible to learners who are as of yet unaware of the contexts differ from those which adult speakers commonly assume license negation.

The issue is, in large part, whether licensing *is* triggering. In other words, if a particular context licenses the use of some linguistic form will this context be sufficient to trigger it as such (i.e., actively accessing NOT in that moment). The force of the present argument is that we can not reverse-engineer competency from access (in the absence of a theory) without potentially drawing arbitrary lines as to what constitutes true negation vs. some ersatz strategy. Indeed, when graphically characterizing an object-tracking model which can explain the initial two cup findings (Cesana-Arlottii et al., 2018), Jasbi (2018) includes a relatively subtle "deletion" operation comparable to the one provided for a syllogistic model which receives a large red X. No time is spent motivating why deletion in one case is constitutive of negation but not the other, and instead the presence of deletion in both suggests the real issue may lay in the setting up of disjuncts. To argue on the basis of potential alternative strategies requires not just that 1) said strategy could be used instead of negation and 2) it is used in the tasks which demonstrate success at earlier ages, but to additionally argue that there is evidence against the use of negation. The only argument to this effect in any of the papers is that there is evidence of object-tracking in other domains earlier, but extending the logic of Morgan's Canon[23] such that earlier available cognitive descriptions are prima facie prioritized in addition to behavioral descriptions risks privileging certain conceptions of negation. This is especially complicated by the fact that object-tracking's early arrival and domain-generality certainly play a role in making certain construals more or less accessible (e.g., any construal which requires noticing a change over time). It is therefore a totally live possibility that object tracking itself employs negation, or is otherwise implicated in how events are parsed and structured such that negation can be applied correctly (e.g., bound to the right items, has the right scope). While at the level of scientific discourse controlling for alternative strategies is inarguably a useful approach, the position being argued for here is in how one should interpret such suggestions in light of the difficulty in developing a definition, or learning theory thereof. The basic point is that immense care should be taken to ensure that a lack of evidence of some competency in childhood speaks to more than the degree of certainty science currently can guarantee in conceptual matters.

The big remaining question for proponents of Access Discontinuity is what sorts of contexts could reasonably be argued to trigger negation. As argued above and explicitly put by Harman (1986a), because such a theory of access is a psychological theory it cannot guarantee the formal certainty required for competence theories. Therefore, an Access Discontinuity theorist is still welcome to study word learning, observation, reference, and the facts about the contexts (both linguistic and non-) that contribute to acquisition contra some initial despair at Mad-Dog Nativism (Landau, 2000). At the same time, they are discouraged from imputing a discontinuity of competency where a discontinuity of access may similarly explain the results at hand, in addition to being compatible with hypothesis testing and more tractable. Chapter 2

[23] As quoted in Sober, 1998: "In no case may we interpret an action as the outcome of the exercise of a higher psychical faculty, if it can be interpreted as the outcome of the exercise of one which stands lower in the psychological scale."

should therefore be read in this light, as it demonstrates straightforwardly how particular contexts may make negation more accessible to both children and adults. In it, both children (4-7yo) and adults were shown both success and failure versions of simple everyday actions and asked to describe them. They both rarely negated when describing successes but regularly negated when describing failures - at no point, however, is this taken as potential evidence that negation reduces to the ability to recognize failures, nor that failure is the only viable route of semantic access, nor that “failures” exist in some mind independent way and we merely internalize it. Instead, it starts from the premise that it’s likely difficult for children to tell when negation is felicitous, both because of performance limitations (e.g., theory of mind, inhibitory control) and because they simply may not have access to relevant information, perhaps as a consequence of performance limitations (e.g., despite “watching” the weather channel with an adult, an infant would be unable to impute any beliefs about the weather onto others). Of course, as has been argued, negation is still learned quite early and thus is likely not the result of explicit practice, but rather from normal everyday uses. Given these difficulties in inferring when someone might use negation, Chapter 2 argues that it is therefore worthwhile to investigate potential “null-discourse” contexts in which a thinker who can think of negation, but does not know yet know how useful it is, might actually be inclined to think of it spontaneously. If a young learner with such a tendency encounters a speaker with a related tendency, then an experience of described failure could serve as a highly informative labelling instance for mapping a negative construal to a negative sentence. Of course, as Study 2 found, there are also alternate construals one may develop in the moment (e.g., “trying to” rather than “failing to”), syntax can further constrain the hypothesis to encourage the mapping of negation. Syntax is likely additionally critical since, presumably, negative construals are also structured somehow, and since individuals are unlikely to guess “not” in isolation (see Chapter One). Mapping NOT onto *not* may not require recapitulating the history of logic,[24] but that's not to say learning *nothing* is easy.

General Discussion

Developmental psychologists have long documented developmental discontinuities be it on behaviors or tasks, and the pressing question is why any patterns at all? Why are children not descriptively equivalent to adults? A popular approach is to take these observations and assume a qualitative change between children and adults, this way lies conceptual discontinuity. Another approach that I have outlined above is to deny that there is any conceptual discontinuity: no qualitative differences between children and adults, at least conceptually speaking. The goal of this paper was to explicate arguments latent in the writings of conceptual continuity theorists, but never united under a catchy name. If this paper has succeeded, there is now a new axis as concerns child conceptual development: how concepts are activated changes over time. I have suggested that there's good reason this must be at least trivially true (e.g., words certainly come to activate concepts). But, whether it explains the effects is a matter of continuing empirical support.

If children do learn concepts, then it is certainly hard work, and so it should be no surprise that learning about learning concepts is similarly hard work. This chapter argued that despite the formal definability of logical concepts, they are no more amenable to concept learning accounts than non-logical concepts. That is, Mad-Dog Nativism does not in fact leave open the possibility that such concepts are learned via their inferential role as some have assumed. This is because the conceptual literature of the past expected definition to do double duty as a learning story, that is, from knowing that BLABBIT = BLUE RABBIT, an adult can find a rabbit in which blue adheres and begin labelling it in front of a child. This is because there was an assumed relationship to reference (via the senses), but the same cannot be assumed for

[24] Well, a particular history of logic which seems to exclude intuitionistic approaches as demonstrated.

logic and so the natural question arises of how one can point at nothing, or falsity, or contrariety, or disjunction, or anything. For any scene, there are a potentially intractably large number of possible logical descriptions, and so the problem of holism (and relatedly semantic ambiguity) resurfaces with a vengeance. Unless one has a theory of access, the problem of learning complex logical concepts is rendered as intractable if not more than learning complex non-logical concepts like DOG. As argued, such a theory is unlikely to be formalizable, and thus unlikely to allow for the direct translation of facts about access to facts about competence.

This paper additionally argues that, as a result of the difficulty even definable contexts pose to concept learning, the literature would do well to rethink analogizing positions in the conceptual sphere to claims about competence and performance. Though such an analogy may have certain pedagogical uses, it obscures that the active site of disagreement is not primarily over whether a performance account can be given, but in whether access constitutes competence. As such, this paper ultimately argues that while Access Discontinuity accounts have been put forth to explain linguistic discontinuities, the same basic principles can be extended to explain discontinuities in non-linguistic behavior. Therefore, the onus is not on Competency Continuity theorists to prove performance limitations explain early failures in negation tasks, but instead on Competency Discontinuity theorists to either 1) argue that their tasks are such that a concept is not just licensed but triggered or 2) develop a formal theory of access. In the absence of either, such research risks being unlikely to make useful contact with word learning research, as it may simply identify cases in which negation is less accessible to children. To be fair, all research, of course, comes with risk, and it's not necessarily a fault of science that it cannot yet guarantee logical certainty. However, there is one final reason an Access Discontinuity is important for cognitive scientists to seriously consider.

The assumption of Access Continuity is, on the one hand, our best guide to develop experimental conditions under which children can display some conceptual ability, and on the other hand, a consequence about metaphysical commitments about what is in the world as opposed to the mind. It is often the difference between chalking up a failure to distinguish in children as a failure of their conceptual abilities instead of our intuitions about possible routes of conceptual (semantic) access. This form of skepticism forms an important pole in the dialectic of developmental psychology, but it can lead to problematic assumptions about non-children. For example, the Pirahã, whose language lacks discrete number terms, have been noted to fail at numerical tasks (Frank et al, 2008a) whereas speakers of languages with number terms succeed. Carey (2009) and others have suggested that they therefore lack a conceptual ability (i.e., they cannot think of number thoughts). However, if one takes Access Discontinuity seriously, it's just as likely that they do, but not like us. Indeed, as Access Discontinuity would predict, speakers of languages with numbers fail at the very same tasks when incapable of recruiting their linguistic abilities. Whether discontinuity theorists intend to or not, the rhetoric of conceptual achievement may serve as justification for differential treatment that eventually leads to the exploitation and deprivation of a community. So goes the story of IQ testing which was created in good will as an educational tool (Nicolas et al, 2013), but is now mostly used to justify the superiority of some over others (Herrnstein & Murray, 2010; for example of such rhetoric in research see Clark et al, 2020: RETRACTED). Given the historical treatment of indigenous communities globally as well as domestically, we may want to ask not just what can be argued, but how those arguments may be used and how they may stick around.

APPENDIX A: SUPPLEMENTAL MATERIALS FOR CHAPTER ONE

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.

APPENDIX B: SUPPLEMENTAL STATISTICS FOR CHAPTER ONE

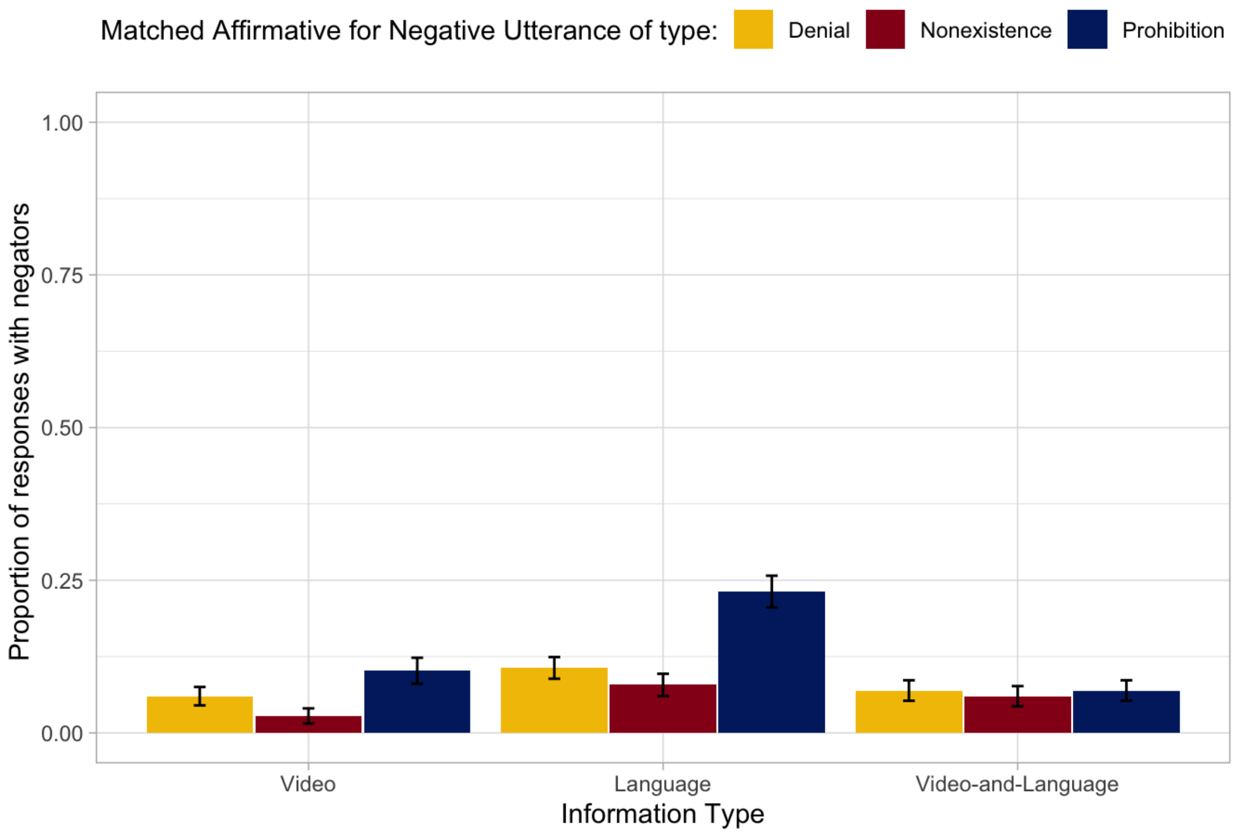

Figure S1: Proportion of negators included in participant responses for matched affirmative vignettes by the Negation Function that each affirmative is paired to. Data are presented by Negation Function (Denial, Nonexistence, Prohibition) and Information Type (Video Only, Video + Language, Language Only), 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).

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 titled "Sampling Procedure for Stimuli.” 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 ($\beta=1.33$, $z=8.67$, $p<0.001$).

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.

APPENDIX C: LINGUISTIC STIMULI FOR CHAPTER THREE

Idealized Sentence	Study 1a: Content Words	Study 1b: Syntax	Study 1c: Both
One of the girls can not tou toes.	girls, her, one, toes, touch	Zum of the prings can ___ gribe ler prents.	One of the girls can ___ touch her toes.
A woman did not pick up a	pen, pick, woman	A tive did ___ swock up a kurg	A woman did ___ pick up a pen
A person can not open the	door, open, person	A fendle can ___ blark the rurf	A person can ___ open the door
The lamp would not turn on	lamp, turn	The crolt would ___ slont on	The lamp would ___ turn on
A girl can not catch a ball.	ball, catch, girl	A pring can ___ zeb a dax	A girl can ___ catch a ball
A girl can not crush one of t cups.	crush, cup, girl, one	A pring can ___ blick zum of the scosts.	A girl can ___ crush one of the cups.
A girl can not get the plug ir outlet.	fitting, girl, outlet, plug	A pring can ___ rall the trunn into the fintzer.	A girl can ___ get the plug into the outlet.
The toilet will not flush.	flush, toilet	The stace will ___ rudge.	The toilet will ___ flush.
A person can not light a ma	light, match, person	A fendle can ___ glizz a ruld	A person can ___ light a match
The pepper is not coming o	coming, pepper	The florb is ___ blaxing out	The pepper is ___ coming out
The pen is not working.	pen, working	The kurg is ___ siffing.	The pen is ___ working.
A person is not cutting a lea	cutting, leaf, person	A fendle is ___ zubbing a flib	A person is ___ cutting a leaf

Table S3: List of linguistic stimuli used in critical trials. See Methods of relevant study in Chapter 2 for additional detail.

APPENDIX D: ADDITIONAL STATISTICS FOR CHAPTER THREE

Alternate Constructions

Employing an explicit negator is not the only way to describe a failure. Instead, one could describe what someone is failing at as an attempt to do something. For example, participants regularly described the actors in the scene as “trying to” complete some actions (hereafter referred to as Try responses; e.g., “The person is trying to open the door.”). In addition to Try responses, there is also the issue raised by terms which are argued to be implicitly negative (e.g., “The girls try to touch their toes but one failed her goal.”; see Dähl, 2010). To investigate both these potentials, we coded participant responses to indicate whether they included either Try constructions (*try, trying, tried, tries*) or implicit negators (*forbid, forbear, avoid, prohibit, stop, fail, forget, doubt, deny*). See Table 1 for a list of proportions of Try and implicit negation responses across Studies 1a-c from those responses which did not already contain explicit negation.

Experiment	Information Condition	Try (Success)	Try (Failure)	Implicit (Success)	Implicit (Failure)
1a: Content Words	Language-Only	0.01	0.04	0	0
1a: Content Words	Video-Only	0.06	0.24	0	0.01
1a: Content Words	Video w. Language	0.03	0.15	0	0
1b: Syntax	Language-Only	0.01	0	0	0
1b: Syntax	Video-Only	0.04	0.18	0	0
1b: Syntax	Video w. Language	0	0.01	0	0
1c: Full Ling.	Language-Only	0	0	0	0
1c: Full Ling.	Video-Only	0.01	0.28	0	0
1c: Full Ling.	Video w. Language	0	0.01	0	0

Table S4: Proportion of responses that included *try* or an implicit negator for Studies 1a-c.

Experiment	Order of Exposures	Try (Success)	Try (Failure)	Implicit (Success)	Implicit (Failure)
2a: Children	Success First	0.095	38.46	0	0
2a: Children	Failures First	0.55	50	0	0.9
2b: Adults	Success First	0.18	31.81	0	3.64
2b: Adults	Failures First	0.64	35.45	0	1.81

Table S5: Proportion of responses that included *try* or an implicit negator for Studies 2-ab.

Mean Length of Child Responses

The following table provides the mean length of responses by participants for both Success and Failure videos. Only complete words were counted (e.g., “uh” and partial words like “th-” would not count). A linear model predicting MLR with Event Type and Exposure Order as predictor found main effects of both (β=-1.419, SE=0.223, z=-6.371 p<0.001; β=-0.684, SE=0.2227 z=-3.07 p<0.005 respectively) but no significant interaction.

Subject	Success/Failure	Difference
Failures First condition		
f3	7.636364/14.454545	-6.8181818
f5	15.272727/16.818182	-1.5454545
f7	8.363636/11.7	-3.3363636
f9	6.909091/10.909091	-4
f14	4.9/6.888889	-1.9888889
f16	7/10	-3
f20	4.636364/5.181818	-0.5454545

f22	9.090909/13.090909	-4
f24	11.454545/16.3	-4.8454545
f25	8.272727/8.272727	0
Successes First condition		
f1	6.100000/7.777778	-1.6777778
f4	8.363636/9.090909	-0.7272727
f6	9.363636/15.181818	-5.8181818
f8	9.454545/14.818182	-5.3636364
f10	6.8/7.888889	-1.0888889
f13	6.7/7.363636	-0.6636364
f15	9.8/12.363636	-2.5636364
f17	3.181818/4.200000	-1.0181818
f19	6.9/8.090909	-1.1909091
f21	5.272727/10.6	-5.3272727

Table S6: Mean length of responses given by the child participants in Study 2 given by Order of Exposure.

BIBLIOGRAPHY

Anscombe, G.E.M. An Introduction to Wittgenstein's Tractatus. South Bend, Ind: St. Augustine's Press, 1967.

Aristotle. Categories and De Interpretatione. Translated by J.L. Ackrill. Clarendon Press, 1975.

Armstrong, Sharon Lee, Lila R Gleitman, and Henry Gleitman. "What Some Concepts Might Not Be." Cognition 13, no. 3 (1983): 263–308.

Auden, W.H. Another Time. London: Faber & Faber limited, 1940.

Austin, Keith, Anna Theakston, Elena Lieven, and Michael Tomasello. "Young Children's Understanding of Denial." Developmental Psychology 50, no. 8 (2017): 2061. https://doi.org/10.1037/a0037179.

Axel Arturo Barceló, Aspeitia. "What Does '&' Mean?" The Proceedings of the Twenty-First World Congress of Philosophy 5 (2007): 45–50.

Bacon, Francis, Lisa Jardine, and Michael Silverthorne. The New Organon. Cambridge Texts in the History of Philosophy. Cambridge [U.K.] ; New York: Cambridge University Press, 2000.

Beckett, Samuel. "Worstward Ho." In Nohow On, 1989.

Berkeley, George. "The Principles of Human Knowledge." Edited by Jonathan Bennett, 1710.

Bermúdez, José Luis. Thinking without Words. Oxford University Press, 2007.

Biederman, Irving, Robert J. Mezzanotte, and Jan C. Rabinowitz. "Scene Perception: Detecting and Judging Objects Undergoing Relational Violations." Cognitive Psychology 14, no. 2 (April 1982): 143–77. https://doi.org/10.1016/0010-0285(82)90007-X.

Bloom, Lois. Language Development; Form and Function in Emerging Grammars. MIT Research Monograph, No. 59. Cambridge, Mass: M.I.T. Press, 1970.

Bohn, Manuel, Josep Call, and Christoph J. Völter. “Evolutionary Precursors of Negation in Non-Human Reasoning.” In The Oxford Handbook of Negation, by Manuel Bohn, Josep Call, and Christoph J. Völter, 576–88. edited by Viviane Déprez and M. Teresa Espinal. Oxford University Press, 2020. https://doi.org/10.1093/oxfordhb/9780198830528.013.40.

Brouwer, L.E.J. “On the Significance of the Principle of Excluded Middle in Mathematics, Especially in Function Theory.” In From Frege to Gödel, A Source Book in Mathematical Logic, by J van Heijenoort, 1977.

Cameron-Faulkner, Thea, Elena Lieven, and Anna Theakston. “What Part of No Do Children Not Understand? A Usage-Based Account of Multiword Negation.” Journal of Child Language 34, no. 2 (May 2007): 251–82. https://doi.org/10.1017/S0305000906007884.

Carey, Susan. “On Learning New Primitives in the Language of Thought: Reply to Rey: On Learning New Primitives in the Language of Thought: Reply to Rey.” Mind & Language 29, no. 2 (March 2014): 133–66. https://doi.org/10.1111/mila.12045.

———. The Origin of Concepts. Oxford Series in Cognitive Development. Oxford ; New York: Oxford University Press, 2009.

———. “Why Theories of Concepts Should Not Ignore the Problem of Acquisition.” Disputatio 7, no. 41 (November 1, 2015): 113–63. https://doi.org/10.2478/disp-2015-0008.

Carroll, Lewis. Through the Looking-Glass and What Alice Found There. Philadelphia: Henry Altemus Company, 1872.

Carvalho, Alex, Cécile Crimon, Axel Barrault, John Trueswell, and Anne Christophe. “‘ 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, February 12, 2021. https://doi.org/10.1111/desc.13085.

Cesana-Arlotti, Nicoló, and Justin Halberda. “A Continuity in Logical Development: Domain-General Logical Inference by 2.5-Year-Old Toddlers.” Preprint. PsyArXiv, November 9, 2021. https://doi.org/10.31234/osf.io/qzxkp.

———. “Domain-General Logical Inference by 2.5-Year-Old Toddlers.” Preprint. PsyArXiv, November 9, 2021. https://doi.org/10.31234/osf.io/qzxkp.

Cesana-Arlotti, Nicoló, Ana Martín, Ernő Téglás, Liza Vorobyova, Ryszard Cetnarski, and Luca L. Bonatti. “Precursors of Logical Reasoning in Preverbal Human Infants.” Science 359, no. 6381 (March 16, 2018): 1263–66. https://doi.org/10.1126/science.aao3539.

Choi, Soonja. “The Semantic Development of Negation: A Cross-Linguistic Longitudinal Study.” Journal of Child Language 15, no. 3 (1988): 517–31.

Chomsky, Carol. “The Acquisition of Syntax in Children from 5 to 10.” Cambridge, Mass: MIT, 1969.

Churchland, Patricia Smith. “Fodor on Language Learning.” Synthese 38, no. 1 (May 1978): 149–59. https://doi.org/10.1007/BF00484757.

Churchland, Paul M. “Eliminative Materialism and the Propositional Attitudes.” The Journal of Philosophy 78, no. 2 (February 1981): 67. https://doi.org/10.2307/2025900.

———. “Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor.” Philosophy of Science 55, no. 2 (June 1988): 167–87. https://doi.org/10.1086/289425.

Clark, Cory J., Bo M. Winegard, Jordan Beardslee, Roy F. Baumeister, and Azim F. Shariff. “RETRACTED: Declines in Religiosity Predict Increases in Violent Crime—but Not

Among Countries With Relatively High Average IQ.” Psychological Science 31, no. 2 (February 2020): 170–83. https://doi.org/10.1177/0956797619897915.

Cowie, Fiona. “The Logical Problem of Language Acquisition.” Synthese 111, no. 1 (1997).

Cuccio, Valentina. “Is Embodiment All That We Need? Insights from the Acquisition of Negation.” Biolinguistics 6, no. 3–4 (2012): 259–75.

———. “On Negation. What Do We Need to ‘Say No’?” Rivista Italiana Di Filosofia Del Linguaggio, no. 4 (December 30, 2011): 47–55. https://doi.org/10.4396/20111205.

Dimroth, Christine. “The Acquisition of Negation.” In The Expression of Negation, edited by Laurence R. Horn, 39–72. DE GRUYTER MOUTON, 2010. https://doi.org/10.1515/9783110219302.39.

Dummett, Michael. The Logical Basis of Metaphysics. 2. impression. London: Duckworth, 1995.

Elbourne, Paul D. Meaning: A Slim Guide to Semantics. Oxford Linguistics. Boston, MA ; New York, NY: Oxford University Press, 2011.

Farennikova, Anna. “Seeing Absence.” Philosophical Studies 166, no. 3 (December 2013): 429–54. https://doi.org/10.1007/s11098-012-0045-y.

Feiman, Roman, Shilpa Mody, and Susan Carey. “The Development of Reasoning by Exclusion in Infancy.” Cognitive Psychology 135 (June 2022): 101473. https://doi.org/10.1016/j.cogpsych.2022.101473.

Feiman, Roman, Shilpa Mody, Sophia Sanborn, and Susan Carey. “What Do You Mean, No? Toddlers’ Comprehension of Logical ‘No’ and ‘Not.’” Language Learning and Development 13, no. 4 (October 2, 2017): 430–50. https://doi.org/10.1080/15475441.2017.1317253.

Fitch, Allison, Sudha Arunachalam, and Amy M. Lieberman. “Mapping Word to World in ASL: Evidence from a Human Simulation Paradigm.” Cognitive Science 45, no. 12 (December 2021). https://doi.org/10.1111/cogs.13061.

Fodor, J.A., M.F. Garrett, E.C.T. Walker, and C.H. Parkes. “Against Definitions.” Cognition 8, no. 3 (January 1980): 263–367. https://doi.org/10.1016/0010-0277(80)90008-6.

Fodor, Jerry. Concepts: Where Cognitive Science Went Wrong. Oxford University Press, 1988.

———. “Having Concepts: A Brief Refutation of the Twentieth Century.” Mind and Language 19, no. 1 (February 2004): 29–47. https://doi.org/10.1111/j.1468-0017.2004.00245.x.

———. “On the Impossibility of Acquiring ‘More Powerful’ Structures.” In Language and Learning: The Debate between Jean Piaget and Noam Chomsky, 143–62. Cambridge, Massachusetts: Harvard University Press, 1980.

Fodor, Jerry A. “A Reply to Churchland’s ‘Perceptual Plasticity and Theoretical Neutrality.’” Philosophy of Science 55, no. 2 (June 1988): 188–98. https://doi.org/10.1086/289426.

———. Concepts: Where Cognitive Science Went Wrong. Oxford Cognitive Science Series. Oxford : New York: Clarendon Press ; Oxford University Press, 1998.

———. Psychosemantics: The Problem of Meaning in the Philosophy of Mind. 4. print. Explorations in Cognitive Science 2. Cambridge, MA: MIT Press, 1998.

———. The Language of Thought. The Language & Thought Series. New York: Crowell, 1975.

Fodor, Jerry, and Ernie Lepore. “All at Sea in Semantic Space: Churchland on Meaning Similarity.” The Journal of Philosophy 96, no. 8 (1999): 381–403.

———. “Analyticity Again.” In The Blackwell Guide to the Philosophy of Language, edited by Michael Devitt and Richard Hanley, 1st ed., 114–30. Wiley, 2006. https://doi.org/10.1002/9780470757031.ch6.

Fox, John, and Sanford Weisberg. “Multivariate Linear Models in R.” Los Angeles: Thousand Oaks, 2011.

Franck, Julie. “Syntactic Encoding: Novel Insights Into the Relationship Between Grammar and Processing.” In The Handbook of Psycholinguistics, edited by Eva M. Fernández and Helen Smith Cairns, 1st ed., 13–37. Wiley, 2017. https://doi.org/10.1002/9781118829516.ch2.

———. “Syntactic Encoding: Novel Insights Into the Relationship Between Grammar and Processing.” In The Handbook of Psycholinguistics, edited by Eva M. Fernández and Helen Smith Cairns, 1st ed., 13–37. Wiley, 2017. https://doi.org/10.1002/9781118829516.ch2.

Frank, Michael C., Daniel L. Everett, Evelina Fedorenko, and Edward Gibson. “Number as a Cognitive Technology: Evidence from Pirahã Language and Cognition.” Cognition 108, no. 3 (September 2008): 819–24. https://doi.org/10.1016/j.cognition.2008.04.007.

Frege, Gottlob. The Foundations of Arithmetic, 1884.

Friston, Karl. “The Free-Energy Principle: A Unified Brain Theory?” Nature Reviews Neuroscience 11, no. 2 (February 2010): 127–38. https://doi.org/10.1038/nrn2787.

Friston, Karl, and Stefan Kiebel. “Predictive Coding under the Free-Energy Principle.” Philosophical Transactions of the Royal Society B: Biological Sciences 364, no. 1521 (May 12, 2009): 1211–21. https://doi.org/10.1098/rstb.2008.0300.

Gabbay, Dob. “What Is Negation in a System?” Studies in Logic and the Foundations of Mathematics 124 (1987): 95–112.

Gabbay, Dov M., and Anthony Hunter. “Negation and Contradiction.” In What Is Negation?, edited by Dov M. Gabbay and Heinrich Wansing, Vol. 13. Applied Logic Series. Dordrecht: Springer Netherlands, 1999. https://doi.org/10.1007/978-94-015-9309-0.

Gabbay, Dov M., and Heinrich Wansing, eds. What Is Negation? Vol. 13. Applied Logic Series. Dordrecht: Springer Netherlands, 1999. https://doi.org/10.1007/978-94-015-9309-0.

Geach, Peter Thomas. Logic Matters. 1. California paperback print. Berkeley, Calif.: Univ. of California Pr, 1980.

Gentner, Dedre. “On Relational Meaning: The Acquisition of Verb Meaning.” Child Development 49 (1978): 988–98.

Gillette, Jane, Henry Gleitman, Lila Gleitman, and Anne Lederer. “Human Simulations of Vocabulary Learning.” Cognition 73, no. 2 (December 1999): 135–76. https://doi.org/10.1016/S0010-0277(99)00036-0.

Gleitman, Lila. “The Structural Sources of Verb Meanings.” Language Acquisition 1, no. 1 (1990): 3–55.

Gleitman, Lila, and Henry Gleitman. “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, no. 1 (1992): 31–35.

Gleitman, Lila R., Kimberly Cassidy, Rebecca Nappa, Anna Papafragou, and John C. Trueswell. “Hard Words.” Language Learning and Development 1, no. 1 (January 2005): 23–64. https://doi.org/10.1207/s15473341lld0101_4.

Gleitman, Lila R., and John C. Trueswell. “Easy Words: Reference Resolution in a Malevolent Referent World.” Topics in Cognitive Science 12, no. 1 (January 2020): 22–47. https://doi.org/10.1111/tops.12352.

Goldin-Meadow, Susan, Susan C. Levine, Larry V. Hedges, Janellen Huttenlocher, Stephen W. Raudenbush, and Steven L. Small. “New Evidence about Language and Cognitive Development Based on a Longitudinal Study: Hypotheses for Intervention.” American Psychologist 69, no. 6 (2014): 588–99. https://doi.org/10.1037/a0036886.

Gomes, Victor, Rebecca Doherty, Daniel Smits, Susan Goldin-Meadow, John C. Trueswell, and Roman Feiman. “It’s Not Just What We Don’t Know: The Mapping Problem in the Acquisition of Negation.” Cognitive Psychology 145 (September 2023): 101592. https://doi.org/10.1016/j.cogpsych.2023.101592.

Gomes, Victor, Yubin Huh, and John Trueswell. “Not What You Expect: The Relationship between Violation of Expectation and Negation.” Proceedings of the 42nd Annual Conference of the Cognitive Science Society, 2020. https://doi.org/10.17605/OSF.IO/MD832.

Hafri, Alon, and Chaz Firestone. “The Perception of Relations.” Trends in Cognitive Sciences 25, no. 6 (2021): 492.

Hafri, Alon, Anna Papafragou, and John C. Trueswell. “Getting the Gist of Events: Recognition of Two-Participant Actions from Brief Displays.” Journal of Experimental Psychology: General 142, no. 3 (2013): 880–905. https://doi.org/10.1037/a0030045.

Hafri, Alon, John C. Trueswell, and Brent Strickland. “Encoding of Event Roles from Visual Scenes Is Rapid, Spontaneous, and Interacts with Higher-Level Visual Processing.” Cognition 175 (June 2018): 36–52. https://doi.org/10.1016/j.cognition.2018.02.011.

Halberda, Justin. “Perceptual Input Is Not Conceptual Content.” Trends in Cognitive Sciences 23, no. 8 (August 2019): 636–38. https://doi.org/10.1016/j.tics.2019.05.007.

Hand, Michael. “Antirealism and Falsity.” In What Is Negation?, edited by Dov M. Gabbay and Heinrich Wansing, Vol. 13. Applied Logic Series. Dordrecht: Springer Netherlands, 1999. https://doi.org/10.1007/978-94-015-9309-0.

Harman, Gilbert. Change in View: Principles of Reasoning. Cambridge, Mass: MIT Press, 1986.

———. “The Meaning of Logical Constants.” In Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson, edited by Ernest LePore. Oxford, UK ; New York, NY, USA: Blackwell, 1986.

Herrnstein, Richard J., and Charles Murray. The Bell Curve: Intelligence and Class Structure in American Life. New York: Free Press, 1996.

Heyting, Arend. Intuitionism: An Introduction. 3rd ed. Studies in Logic and the Foundations of Mathematics. Amsterdam New York Oxford: North-Holland, 1980.

Hochmann, Jean-Rémy. “Cognitive Precursors of Negation in Preverbal Infants.” In The Oxford Handbook of Negation, by Jean-Rémy Hochmann, 588–98. edited by Viviane Déprez and M. Teresa Espinal. Oxford University Press, 2020. https://doi.org/10.1093/oxfordhb/9780198830528.013.39.

Hochmann, Jean-Rémy, and Juan M. Toro. “Negative Mental Representations in Infancy.” Cognition 213 (August 2021): 104599. https://doi.org/10.1016/j.cognition.2021.104599.

Horn, Laurence R. A Natural History of Negation. The David Hume Series. Stanford, Calif: CSLI, 2001.

———, ed. “The Acquisition of Negation.” In The Expression of Negation. Berlin, New York: DE GRUYTER MOUTON, 2010. https://doi.org/10.1515/9783110219302.39.

Hummer, Peter, Heinz Wimmer, and Gertraud Antes. “On the Origins of Denial Negation.” Journal of Child Language 20, no. 3 (October 1993): 607–18. https://doi.org/10.1017/S0305000900008503.

Jasbi, Masoud, Manuel Bohn, Bria Long, Abdellah Fourtassi, David Barner, and Michael C Frank. “Comment on Cesana-Arlotti et al. (2018).” PsyArXiv, 2019.

Ji, Yue, and Anna Papafragou. “Boundedness in Event Cognition: Viewers Spontaneously Represent the Temporal Texture of Events.” Journal of Memory and Language 127 (December 2022): 104353. https://doi.org/10.1016/j.jml.2022.104353.

Johnson-Laird, P. N., and J. M. Tridgell. “When Negation Is Easier than Affirmation.” Quarterly Journal of Experimental Psychology 24, no. 1 (February 1972): 87–91. https://doi.org/10.1080/14640747208400271.

Johnson-Laird, Philip Nicholas. The Computer and the Mind: An Introduction to Cognitive Science. Cambridge, Mass: Harvard Univ. Press, 1988.

Jowlett, Benjamin. The Complete Works of Plato. Cork: E-artnow Editions, 2013.

Keller, Georg B., and Thomas D. Mrsic-Flogel. “Predictive Processing: A Canonical Cortical Computation.” Neuron 100, no. 2 (October 2018): 424–35. https://doi.org/10.1016/j.neuron.2018.10.003.

Kessel, Frank S. “The Role of Syntax in Children’s Comprehension from Ages Six to Twelve.” Monographs of the Society for Research in Child Development 35, no. 6 (September 1970): iii. https://doi.org/10.2307/1165662.

Kim, Kyung J. “Development of the Concept of Truth-Functional Negation.” Developmental Psychology 21, no. 3 (1985).

Klimacka, Lindsay, and Karen Brunger. “Elicitation in Verb Morphology.” Child Language Teaching and Therapy 15, no. 3 (1999): 247–59.

Landau, Barbara. “Concepts, the Lexicon and Acquisition: Fodor’s New Challenge.” Mind and Language 15, no. 2 & 3 (April 2000): 319–26. https://doi.org/10.1111/1468-0017.00136.

Lasnik, Howard. “Analyses of Negation in English.Pdf.” Doctoral Dissertation. Massachusetts Institute of Technology, 1972.

Lenzen, Wolfgang. “Necessary Conditions for Negation Operators.” In Negation: A Notion in Focus, edited by Heinrich Wansing. Perspektiven Der Analytischen Philosophie 7. Berlin: de Gruyter, 1996.

Locke, John. An Essay Concerning Human Understanding. London: Thomas Basset, 1690.

Malt, Barbara C. “Features and Beliefs in the Mental Representation of Categories.” Journal of Memory and Language 29, no. 3 (June 1990): 289–315. https://doi.org/10.1016/0749-596X(90)90002-H.

Margolis, Eric, and Stephen Laurence, eds. Concepts: Core Readings. Cambridge, Mass: MIT Press, 1999.

McDermott-Hinman, Annika, and Roman Feiman. “The Development of Negation in Language and Thought.” In Perspectives on Negation: Views from across the Language Sciences, in press. https://bpb-us-w2.wpmucdn.com/sites.brown.edu/dist/8/25/files/2024/02/McDermott-Hinman-Feiman-negation-chapter-final-in-press-1c50f65bf780c2c8.pdf.

McNeill, D., and N. McNeill. “A Question in Semantic Development - What Does a Child Mean When He Says ‘No.’” Ann Arbor, MI: Michigan University, Center for Research on Language and Behavior, 1967.

Miller, George A., and Philip N. Johnson-Laird. Language and Perception. Reprint 2014. Cambridge, MA: Harvard University Press, 2013.

Mody, Shilpa, and Susan Carey. “The Emergence of Reasoning by the Disjunctive Syllogism in Early Childhood.” Cognition 154 (September 2016): 40–48. https://doi.org/10.1016/j.cognition.2016.05.012.

Murphy, Gregory. “What Are Categories and Concepts.” In The Making of Human Concepts, edited by D Mareschal, P Quinn, and S.E.G. Lea, 11–28. Oxford: Oxford University Press, 2010.

Naigles, Letitia. “Children Use Syntax to Learn Verb Meanings.” Journal of Child Language 17, no. 2 (June 1990): 357–74. https://doi.org/10.1017/S0305000900013817.

Nappa, Rebecca, Allison Wessel, Katherine L. McEldoon, Lila R. Gleitman, and John C. Trueswell. “Use of Speaker’s Gaze and Syntax in Verb Learning.” Language Learning and Development 5, no. 4 (September 17, 2009): 203–34. https://doi.org/10.1080/15475440903167528.

Nelson, David. “A Complete Negationless System.” Studia Logica 32, no. 1 (December 1973): 41–47. https://doi.org/10.1007/BF02123809.

———. “Constructible Falsity.” Journal of Symbolic Logic 14, no. 1 (May 16, 1949): 16–26. https://doi.org/10.2307/2268973.

Nordmeyer, Ann E, and Michael C Frank. “Individual Variation in Children’s Early Production of Negation.” Annual Meeting of the Cognitive Science Society, 2018.

———. “Negation Is Only Hard to Process When It Is Pragmatically Infelicitous.” N Proceedings of the 37th Annual Meeting of the Cognitive Science Society, 2015, 19–25.

Nordmeyer, Ann E., and Michael C. Frank. “Pragmatic Felicity Facilitates the Production and Comprehension of Negation.” Collabra: Psychology 9, no. 1 (January 26, 2023): 67931. https://doi.org/10.1525/collabra.67931.

———. “The Role of Context in Young Children’s Comprehension of Negation.” Journal of Memory and Language 77 (November 2014): 25–39. https://doi.org/10.1016/j.jml.2014.08.002.

Nordmeyer, Ann E, and Jill G de Villiers. “To Be or Not to Be: Examining the Role of Language in a Concept of Negation.” Philosophy: Faculty Publications, 2019. https://scholarworks.smith.edu/phi_facpubs/65/.

Ockham, William. Part 1 of the Summa Logicae. Translated by Michael Loux, 1349.

Oliva, Aude, and Antonio Torralba. “Modeling the Shape of the Scene: A Holistic Representation of the Spatial Envelope.” International Journal of Computer Vision 42 (2001): 145–75.

Onishi, Kristine H., and Renée Baillargeon. “Do 15-Month-Old Infants Understand False Beliefs?” Science 308, no. 5719 (April 8, 2005): 255–58. https://doi.org/10.1126/science.1107621.

Onishi, Kristine H., Renée Baillargeon, and Alan M. Leslie. “15-Month-Old Infants Detect Violations in Pretend Scenarios.” Acta Psychologica 124, no. 1 (January 2007): 106–28. https://doi.org/10.1016/j.actpsy.2006.09.009.

Papafragou, Anna, Kimberly Cassidy, and Lila Gleitman. “When We Think about Thinking: The Acquisition of Belief Verbs.” Cognition 105, no. 1 (October 2007): 125–65. https://doi.org/10.1016/j.cognition.2006.09.008.

Pea, Roy. “The Development of Negation in Early Child Language.” Doctoral Dissertation. University of Oxford, 1978.

Peacocke, Christopher. A Study of Concepts. 2. MIT Press paperback ed. Representation and Mind. Cambridge, Mass.: MIT Press, 1999.

Piccin, Thomas B., and Sandra R. Waxman. “Why Nouns Trump Verbs in Word Learning: New Evidence from Children and Adults in the Human Simulation Paradigm.” Language Learning and Development 3, no. 4 (August 30, 2007): 295–323. https://doi.org/10.1080/15475440701377535.

Quine, Willard Van Orman. Word and Object. Cambridge, (Mass.): Massachusetts institut of technology press, 1992.

Restall, Greg. “Combining Possibilities and Negations.” Studia Logica 59 (1997).

Reuter, Tracy, Roman Feiman, and Jesse Snedeker. “Getting to No: Pragmatic and Semantic Factors in Two- and Three-Year-Olds’ Understanding of Negation.” Child Development 89, no. 4 (July 2018): e364–81. https://doi.org/10.1111/cdev.12858.

Rey, Georges. “Innate and Learned: Carey, Mad Dog Nativism, and the Poverty of Stimuli and Analogies (Yet Again): Innate and Learned.” Mind & Language 29, no. 2 (March 2014): 109–32. https://doi.org/10.1111/mila.12044.

Rolfs, Martin, Michael Dambacher, and Patrick Cavanagh. “Visual Adaptation of the Perception of Causality.” Current Biology 23, no. 3 (February 2013): 250–54. https://doi.org/10.1016/j.cub.2012.12.017.

Russell, Bertrand. The Metaphysician’s Nightmare. London: Allen & Unwin, 1954.

Sartre, Jean-Paul, and Hazel Estella Barnes. Being and Nothingness: A Phenomenological Essay on Ontology. London: Routledge Classics, 2003.

Seyfarth, Robert, and Dorothy Cheney. “The Assessment by Vervet Monkeys of Their Own and Another Species’ Alarm Calls.” Animal Behaviour 40, no. 4 (October 1990): 754–64. https://doi.org/10.1016/S0003-3472(05)80704-3.

———. “The Assessment by Vervet Monkeys of Their Own and Another Species’ Alarm Calls.” Animal Behaviour 40, no. 4 (October 1990): 754–64. https://doi.org/10.1016/S0003-3472(05)80704-3.

Shea, Nicholas. “New Concepts Can Be Learned: Susan Carey, The Origin of Concepts, Oxford University Press, Oxford, 2009.” Biology & Philosophy 26, no. 1 (January 2011): 129–39. https://doi.org/10.1007/s10539-009-9187-5.

Snedeker, Jesse, Joy Geren, and Carissa L. Shafto. “Starting Over: International Adoption as a Natural Experiment in Language Development.” Psychological Science 18, no. 1 (January 2007): 79–87. https://doi.org/10.1111/j.1467-9280.2007.01852.x.

Sober, E. “Morgan’s Canon.” In The Evolution of Mind, by D.D. Cummins and C Allen, 224–42, 1998.

Sorensen, R. Seeing Dark Things: The Philosophy of Shadows. Oxford: Oxford University Press, 2008.

Soteriou, Matthew. “The Perception of Absence, Space, and Time.” In Perception, Causation, and Objectivity. Oxford University Press, 2011.

Stevens, Jon Scott, Lila R. Gleitman, John C. Trueswell, and Charles Yang. “The Pursuit of Word Meanings.” Cognitive Science 41 (April 2017): 638–76. https://doi.org/10.1111/cogs.12416.

Sullivan, Jessica, Michelle Mei, Andrew Perfors, Erica Wojcik, and Michael C. Frank. “SAYCam: A Large, Longitudinal Audiovisual Dataset Recorded From the Infant’s Perspective.” Open Mind 5 (May 26, 2021): 20–29. https://doi.org/10.1162/opmi_a_00039.

Tennant, Neil. “Negation, Absurdity, and Contrariety.” In What Is Negation?, edited by Dov M. Gabbay and Heinrich Wansing, Vol. 13. Applied Logic Series. Dordrecht: Springer Netherlands, 1999. https://doi.org/10.1007/978-94-015-9309-0.

Thornton, Rosalind. “Elicited Production of Past Tense Affirmative and Negative Sentences in Young Children.” Proceedings of the 42nd Annual Boston University Conference on Language Development, 2017, 764–76.

Thornton, Rosalind, and Kelly Rombough. “The Syntax-PF Interface in Children’s Negative Sentences.” Language Acquisition 22, no. 2 (April 3, 2015): 132–57. https://doi.org/10.1080/10489223.2014.943901.

Thornton, Rosalind, Kelly Rombough, Jasmine Martin, and Linda Orton. “Negative Sentences in Children with Specific Language Impairment.” First Language 36, no. 3 (June 2016): 228–64. https://doi.org/10.1177/0142723716640187.

Trueswell, John C., Tamara Nicol Medina, Alon Hafri, and Lila R. Gleitman. “Propose but Verify: Fast Mapping Meets Cross-Situational Word Learning.” Cognitive Psychology 66, no. 1 (February 2013): 126–56. https://doi.org/10.1016/j.cogpsych.2012.10.001.

Valtonen, Pasi. “The Meaning of Absurdity.” Logique et Analyse, no. 0 (2017): 375–88. https://doi.org/10.2143/LEA.239.0.3237160.

Van Buren, Benjamin, Stefan Uddenberg, and Brian J. Scholl. “The Automaticity of Perceiving Animacy: Goal-Directed Motion in Simple Shapes Influences Visuomotor Behavior Even

When Task-Irrelevant." Psychonomic Bulletin & Review 23, no. 3 (June 2016): 797–802. https://doi.org/10.3758/s13423-015-0966-5.

Villiers, Jill G. De, and Helen B. Tager Flusberg. "Some Facts One Simply Cannot Deny." Journal of Child Language 2, no. 2 (November 1975): 279–86. https://doi.org/10.1017/S0305000900001100.

Wansing, Heinrich. "Negation." In The Blackwell Guide to Philosophical Logic, edited by Lou Goble. Blackwell Philosophy Guides 4. Malden, Mass: Blackwell Publishers, 2001.

———, ed. Negation: A Notion in Focus. Perspektiven Der Analytischen Philosophie 7. Berlin: de Gruyter, 1996.

———. "Negation as Falsity, a Reply to Tennant." In What Is Negation?, edited by Dov M. Gabbay and Heinrich Wansing, Vol. 13. Applied Logic Series. Dordrecht: Springer Netherlands, 1999. https://doi.org/10.1007/978-94-015-9309-0.

Wason, P C. "In Real Life Negatives Are False." Logique et Analyse, 1972, 17–38.

Watson, John S., Gyorgy Gergely, Vilmos Csanyi, Jozsef Topal, Marta Gacsi, and Zsuzsanna Sarkozi. "Distinguishing Logic from Association in the Solution of an Invisible Displacement Task by Children (Homo Sapiens) and Dogs (Canis Familiaris): Using Negation of Disjunction." Journal of Comparative Psychology 115, no. 3 (2001): 219–26. https://doi.org/10.1037/0735-7036.115.3.219.

Woodard, Kristina, Lila R. Gleitman, and John C. Trueswell. "Two- and Three-Year-Olds Track a Single Meaning During Word Learning: Evidence for Propose-but-Verify." Language Learning and Development 12, no. 3 (July 2, 2016): 252–61. https://doi.org/10.1080/15475441.2016.1140581.

Yu, Chen. “A Statistical Associative Account of Vocabulary Growth in Early Word Learning.” Language Learning and Development 4, no. 1 (January 7, 2008): 32–62. https://doi.org/10.1080/15475440701739353.

Yu, Chen, and Linda B. Smith. “Rapid Word Learning Under Uncertainty via Cross-Situational Statistics.” Psychological Science 18, no. 5 (May 2007): 414–20. https://doi.org/10.1111/j.1467-9280.2007.01915.x.

Yuan, Sylvia, and Cynthia Fisher. “‘Really? She Blicked the Baby?’: Two-Year-Olds Learn Combinatorial Facts About Verbs by Listening.” Psychological Science 20, no. 5 (May 2009): 619–26. https://doi.org/10.1111/j.1467-9280.2009.02341.x.

Yurovsky, Daniel, Linda B. Smith, and Chen Yu. “Statistical Word Learning at Scale: The Baby’s View Is Better.” Developmental Science, March 2013, n/a-n/a. https://doi.org/10.1111/desc.12036.

Zehr, Jérémy, and Florian Schwarz. “PennController for Internet Based Experiments (IBEX),” August 30, 2022. https://doi.org/10.17605/OSF.IO/MD832.

Zhang, Yayun, Daniel Yurovsky, and Chen Yu. “Statistical Word Learning Is a Continuous Process: Evidence from the Human Simulation Paradigm.” N Proceedings of the 37th Annual Meeting of the Cognitive Science Society, 2015.