Becoming word meaning experts: Infants' processing of familiar words in the context of typical and
atypical exemplars

Martin Zettersten^{1*}, Haley Weaver^{2*}, & Jenny R. Saffran²

¹Princeton University, Department of Psychology, Princeton, NJ 08544, USA

²University of Wisconsin-Madison, Department of Psychology, 1202 W. Johnson St., Madison, WI
53706, USA

*equal authorship contribution

Stage 1 Registered Report, in-principle acceptance, Child Development

Corresponding author: Martin Zettersten, Department of Psychology, Princeton University,
Princeton, NJ 08544, USA, martincz@princeton.edu

Data availability statement:

The stimuli, experimental materials, and data that support this study are openly available in a publicly accessible repository on the Open Science Framework (OSF) at
https://osf.io/3t8gf/?view_only=d4031533ac8549d5a8682b7819ba7bee.

Acknowledgements

This work was supported by NSF-GRFP DGE-1747503 to MZ and grants from the NICHD awarded to JS (R37HD037466) and the Waisman Center (U54 HD090256).

Word count: 9,089

Abstract

How do infants become word meaning experts? This study will investigate the structure of infants' early lexical representations by manipulating the typicality of exemplars from familiar animal categories. 14- to 18-month-old infants (planned sample: $N=80$; 40 female; expected race/ethnicity: 60% White, 18% Asian, 3% Hispanic, 2% Black, 17% multiple categories) will be tested on their ability to recognize typical and atypical category exemplars after hearing familiar basic-level category labels. The main question is whether infants' word recognition extends to a narrow vs. broad range of category exemplars. The design provides 80% power to detect an effect size of $d=0.32$. Individual differences in word recognition will also be investigated. This study will illuminate how infants' knowledge of word extension develops.

Keywords: word knowledge, semantic knowledge, typicality, categories, vocabulary development, infancy, extension

In their first years of life, infants rapidly learn about the meanings of words. As early as 6 months of age, infants recognize some frequently-heard words as referring to familiar objects (Bergelson & Swingley, 2012; Kartushina & Mayor, 2019; Tincoff & Jusczyk, 1999). This ability expands with age and experience as infants' ability to link familiar words to possible referents becomes more refined (Bergelson, 2020; Bergelson & Aslin, 2017); the number of words for which children can recognize referents increases rapidly across the second year of life (Braginsky et al., 2019; Frank et al., 2021). These word learning skills lay the foundation for children's early language development and predict a range of later linguistic, cognitive, and academic outcomes (Bleses et al., 2016; Marchman et al., 2018).

What is the nature of infants' early word-meaning representations? As a starting point for this question, it is useful to consider the structure of word representations that will eventually emerge over development. Studies designed to uncover adults' knowledge of word meanings reveal rich underlying structure in the semantic associations within lexical networks (e.g., De Deyne et al., 2019; Steyvers & Tenenbaum, 2005) and in the potential referents to which word meanings extend (Goldberg, 2019). Rather than indexing a specific, narrow range of referents, each noun is connected to a rich category of interrelated referents. One fundamental organizing principle in adults' representations of word meanings is typicality. Adults extend the label *dog* to refer both to typical (e.g., beagles) and atypical (e.g., dachshunds) dogs, with typical exemplars and their features accessed more rapidly than atypical exemplars and features (Rosch, 1978). Infants will eventually become word meaning experts who can effortlessly apply the words they know to a broad range of category items - but when and how does this word knowledge emerge?

The development of word meaning representations - broad-to-narrow or narrow-to-broad?

Two general accounts of infants' understanding of the extension of newly-learned words have been advanced in past literature: a "broad-to-narrow" view and a "narrow-to-broad" view (Hirsh-Pasek et al., 2004). On the broad-to-narrow view, infants construe early-learned words as referring to broad

categories of items, subsequently refining their understanding of a word's possible referents based on experience (Waxman & Gelman, 2009; Waxman & Markow, 1995; Yin & Csibra, 2015). This account predicts that infants use initial learning experiences to successfully link words to a wide range of both typical and atypical exemplars early on in development, supported by rich prior knowledge about category structure (Snedeker & Gleitman, 2004) or by general expectations that words will refer to categories (Ferguson & Waxman, 2017; Waxman & Markow, 1995). One source of evidence in support of this view is that children overextend their word meanings to referents that are perceptually similar, categorically related, or thematically associated with the target word referent (Gruendel, 1977; Huttenlocher & Smiley, 1987; Mayor & Plunkett, 2010; Rescorla, 1980). For instance, children may use the word *horse* to correctly refer to horses, but also include other semantically related animal kinds such as goats (Rescorla, 1980). The broad-to-narrow view suggests that as children amass experience with language and how it is used to refer to objects in their environments, they gradually refine their understanding of words' possible referents and become increasingly adept at distinguishing between closely related referents to fine-tune their word meanings. Thus, the broad-to-narrow account predicts that infants' broad early word meanings should have limited sensitivity to idiosyncratic experiences with individual category exemplars. Once infants grow older, individual differences in experience with possible word referents begin to play an increasingly important role in constraining word meanings.

By contrast, the narrow-to-broad view suggests that infants may first form restricted, exemplar-based representations of a word's meaning, and gradually begin to generalize words to encompass broader categories through extended experience with words co-occurring with a variety of exemplars (Ambridge, 2020; Barrett, 1986; Hirsh-Pasek et al., 2004; Smith et al., 2002). This account predicts that infants' early word meanings may be limited to familiar, typical exemplars. Indeed, some accounts of infants' lexical production suggest that they may initially use a label to refer to a specific object prior to correctly generalizing the label to the entire class of items (Gruendel, 1977; Rescorla, 1980). A second prediction is that infants' early word meaning representations should be closely linked to

individual differences in experience, since the specific exemplars that infants encounter while hearing labels (e.g., hearing the word *dog* together with the family pet) will vary substantially and idiosyncratically across infants. In support of this hypothesis, there is evidence that infants extend early words that have highly familiar, salient referents more slowly than other kinds of words. For example, infants are slower to correctly generalize words such as *dog* or *Momma*, perhaps because infants' experience with these particular words is limited to exposure to a specific category referent (i.e., their own mother or pet dog; Rescorla, 1980). In sum, the narrow-to-broad view predicts that infants represent early word meanings as associated with a narrow set of category referents, and gradually incorporate a broader set of items into their word meaning representations through additional experience with words and their referents.

Prior research has provided empirical evidence that supports both accounts and has been limited in its ability to disambiguate these views for several reasons. First, most research investigating how infants generalize labels across category members introduces novel words and artificially constructed categories. While these studies have offered tremendous insight into the dimensions and features that infants use when generalizing word meanings, they leave unclear the nature of infants' early representations of words they encounter frequently in their day-to-day experience. Second, the vast majority of studies investigating familiar word knowledge focus on a limited set of prototypical category referents and rarely contrast infants' ability to link the same word with a range of exemplars (with notable exceptions that will be discussed below, e.g., Garrison et al., 2020; Meints et al., 1999; Perry & Saffran, 2017; Southgate & Meints, 2000). Third, mainly due to limitations in power in infant research, past work has often focused on whether infants exhibit word recognition *across* all participants and items, at the expense of interrogating individual infants' knowledge about specific items, limiting researchers' ability to assess the role of idiosyncratic experience in the development of word meanings. Next, we review past research investigating the development of infants' word meanings, focusing on (a) relationships between lexical and semantic networks, (b) individual

differences in word meanings, and (c) past research on infants' understanding of words in relation to typical and atypical category exemplars.

Structure and variability in children's early representations of word meaning

Lexical-semantic networks. One early window into the nature of infants' word meanings is the development of their early lexical networks. How infants represent the relationships between words in their budding vocabularies yields insight into how infants represent the underlying meanings of words. Past research suggests that infants do not simply learn words as isolated representations; instead, they connect semantically related words from early on in development (Hills et al., 2009; Peters & Borovsky, 2019; Wojcik, 2018; Wojcik & Saffran, 2013). For example, by the age of two, infants distinguish associations between semantically-related words, such as *dog* and *kitty*, from associations between semantically-unrelated words (e.g., *dog* and *juice*) (Willits et al., 2013). The interconnected nature of infants lexical-semantic networks has also been revealed through priming studies (Arias-Trejo & Plunkett, 2010; Mani et al., 2013). For example, infants will orient faster to a target image after hearing a familiar word (e.g., *cookies*) when it is preceded by a semantically related prime word (e.g., *bananas*) compared to an unrelated word (e.g., *house*). Thus, early vocabulary networks are organized around semantic properties of word referents. However, while studies on the development of early vocabulary networks elucidate the semantic relationships between words, they leave open the question of how broadly or narrowly infants represent the meanings of individual words.

Individual differences in word meanings. Children exhibit individual differences in their vocabulary networks, with consequences for how they represent known words and interpret new labels (Beckage et al., 2011; Colunga & Sims, 2017). For example, toddlers who know more terms for categories organized by material (e.g., *wood* or *chalk*) are more likely to generalize the meaning of a novel word based on material (Perry & Samuelson, 2011). The structure of infants' early vocabularies influences how easily they acquire novel words (Borovsky et al., 2016) and how readily they recognize familiar words in the presence of competing referents (Borovsky & Peters, 2019). These types of

individual differences apply not only to children's spontaneous generalization of novel words but point to deeper individual differences in how infants represent the meanings of known words. In one study (Perry & Saffran, 2017), 21-month-olds were presented with familiar objects depicted in typical (e.g., a red strawberry) or atypical colors (e.g., a green strawberry). While infants correctly recognized words in the presence of both typical and atypical target referents, their word processing was significantly less accurate for objects with atypical colors, suggesting that by 21 months, infants are sensitive to prototypical features of familiar referents (like the color of a strawberry) during word recognition. Crucially, the degree to which infants experienced a processing disruption depended on their existing vocabulary structure: children who knew more shape words were less disrupted, perhaps because they had a more strongly shaped-based representation of words such as strawberry. Word meanings likely vary substantially across infants, but this may not be evident when infants are tested only on prototypical category referents.

One recent study illustrates how individual differences in infants' early word meaning representations can be revealed by testing word recognition beyond standardized category exemplars (Garrison et al., 2020). Infants' word recognition for familiar as opposed to unfamiliar prototypical category referents was tested by asking caregivers to supply photographs of objects from infants' home environment. Infants, who ranged in age between 12 and 18 months, were then tested on how quickly and accurately they linked a known label (e.g., *book*) to the caregiver-supplied image of an object familiar to them (e.g., a book commonly read to them in the home) compared to an image of an unfamiliar book. While older infants were similarly able to link both the familiar and the unfamiliar image to the target label, younger infants only showed above-chance word recognition for familiar items - consistent with the hypothesis that infants' early word meanings are associated with a narrow set of (potentially idiosyncratic) familiar exemplars. However, significant differences between familiarity conditions were inconsistent across analyses and no interaction between familiarity and age

was observed, leading the authors to conclude that their data did not provide strong evidence for a narrow-to-broad change in word meanings.

Category member typicality in word recognition. Even young infants can form categories after experiencing perceptually diverse exemplars (Bornstein & Arterberry, 2010; Eimas & Quinn, 1994). Exemplar typicality plays a key role in early category learning; exemplar similarity and prototypicality ease infants' ability to form visual categories (Bauer et al., 1995; Oakes et al., 1997). Moreover, infants' individual experiences with exemplars from a given category influence their perceptual processing of category members (Hurley & Oakes, 2018; Kovack-Lesh et al., 2014) and the degree to which infants develop finer-grained sensitivity to distinctions within basic-level categories (e.g., distinguishing beagles from Saint Bernards within the dog category; Quinn & Tanaka, 2007). Given that infants' developing perceptual categories are organized by exemplar typicality, to what extent are infants' early representations of word meanings also constrained by category typicality?

There is some evidence to suggest that infants initially recognize referents that are prototypical category exemplars more easily than referents that are atypical exemplars when processing familiar words. Meints and colleagues (1999) tested infants' knowledge of known words - ascertained through parental report for each participant - using typical or atypical members of the named category as referents in a looking-while-listening paradigm at 12, 18, and 24 months. Infants were presented with typical or atypical members of the named target category paired with typical or atypical members of a different category, and infants' fixation of the target image was measured after hearing the known label. 12-month-old infants increased their looking to typical category members (e.g., a sparrow) after word onset, but not to atypical category members (e.g., an ostrich). However, older infants (18- and 24-month-olds) recognized both typical and atypical category members in response to a category label. Relatedly, Poulin-Dubois and Sissons (2002) found that 18-month-old infants recognize incomplete images of category members in which some parts of the image were deleted (thus making these items, arguably, atypical exemplars) as instances of familiar words in a preferential looking task. These

findings suggest that infants may initially have relatively narrow representations of known words, and subsequently form broader representations of a word's meaning.

However, limitations in the literature complicate this interpretation, leaving open the question of how word meaning extensions change across development. First, close inspection of Meints et al. (1999)'s analyses suggest a more modest conclusion regarding whether typical exemplars were recognized better than atypical exemplars by younger infants: while significant looking was found for typical exemplars and not for atypical exemplars, no significant difference in looking to typical vs. atypical exemplars was reported for 12-month-olds, and no significant interaction showing change across age was observed. Similarly, Garrison et al. (2020) found suggestive but inconsistent evidence that object familiarity predicted word recognition - while younger infants showed successful word recognition only for familiar items, no consistent evidence supported a difference between familiar and unfamiliar items. Conversely, while some studies find similar patterns of word recognition for typical and atypical exemplars for infants from 18 to 24 months in age (Meints et al., 1999; Robinson, 2002), other results provide evidence for typicality effects among infants within this age window (Perry & Saffran, 2017; Southgate & Meints, 2000). For example, when presented with two exemplars from a given category (one typical and one atypical), 18-month-old and 24-month-old infants preferred to look at the prototypical category member after hearing its label (Southgate & Meints, 2000). Thus, while some studies provide evidence in support of a narrow-to-broad developmental change in infants' early word meanings, questions remain both about whether infants' early word representations differ by typicality and the extent to which typicality effects on infants' word recognition disappear by 18 months.

A second open question in the literature is the role of experience in shaping early word meanings. A key prediction of the narrow-to-broad account is that infants' early word meaning representations should be principally associated with the specific exemplars that are highly representative of infants' idiosyncratic environments. However, there have been few direct tests of the

degree to which infants' word representations vary across category exemplars as a function of experience. One limitation of past studies is that they have typically tested only one typical or atypical referent per category, which may obscure the extent to which infants differ in how they generalize a given word's meaning across a broader range of category members. While one study (Robinson, 2002) did test familiar word recognition for a wider range of category members in 17- to 25-month-olds, the results revealed no evidence of word recognition for either typical or atypical items among older infants, limiting the conclusions that can be drawn. A likely explanation of this absence of word recognition lies in the unusual variation of the preferential looking procedure used in the study, which included an 8-second baseline window and an 8-second looking window post-label onset, far longer than designs used to consistently elicit familiar word recognition (see e.g., Fernald et al., 2008). A second limitation of the current literature is that few studies have sought to directly quantify infants' experience in relation to their word representations (with some notable exceptions: Garrison et al., 2020; Robinson, 2002; Southgate & Meints, 2000). Southgate and Meints (2000) quantified object experience through caregiver report but did not relate this measure to infant behavior. Garrison et al. (2020) found no influence of the degree of similarity between familiar and unfamiliar items on infants' word recognition, but the typicality range of the items used in the study was limited due to the fact that even the unfamiliar items tended to be prototypical category members. Therefore, it currently remains unclear to what extent infants' early word meanings vary across individuals and whether such differences stem from differences in individual experience.

In sum, two key questions remain from past work: how narrow are infants' early word meanings and to what degree do early word meanings vary with experience? First, while some past studies provide evidence that early word meanings are narrowly focused on typical items, the literature reveals inconsistencies that complicate this interpretation. Some authors argue for a narrow-to-broad developmental change such that there is a typicality advantage in early word meanings that dissipates by 18 months (Meints et al., 1999), while other studies suggest that early word meanings may not vary

substantially according to experience typicality (Garrison et al., 2020) or that typicality effects in word recognition continue to be found beyond 18 months (Perry & Saffran, 2017). Second, while a narrow-to-broad account predicts that infants' early word representations should vary based on experience with specific category exemplars, we still have little direct evidence relating experiential factors to variability in the breadth of infants' early word meanings. Conclusions about how broadly children generalize particular words across a given category are limited by the small number of category exemplars used in past studies and the limited number of attempts to relate infant word meanings directly to typicality of prior experience. In the current study, we seek to clarify these two questions by investigating developmental change in infants' word representations tested across a broader set of category exemplars, while using caregiver report to quantify exemplar typicality for each infant.

The present study

Are infants' early word representations initially limited to a narrow set of exemplars (e.g., that the word *dog* refers to a prototypical category member like Labrador) or do these representations include broad sets of exemplars from early in development (e.g., that *dog* may also refer to less representative category members such as chihuahuas, greyhounds, and pugs)? In the current study, we will investigate the variability in infants' lexical representations across development by manipulating the typicality of exemplars from highly familiar categories. Specifically, we will test infants (14-18 months of age) on their ability to recognize typical and atypical category exemplars in a looking-while-listening experiment (Fernald et al., 2008). The words will be a set of basic-level category nouns (*dog*, *bird*, *cat*, *fish*) that most infants comprehend by 10-15 months of age (supplementary materials S2, Table S1; Frank et al., 2017). A small set of categories will be used to ensure that we can include a wider range of exemplars per category, varying in typicality. Infants will be tested in two sessions in order to maximize the number of observations per infant, and caregivers will complete a questionnaire on their infants' experiences with each category. Our study pursues three main aims.

Aim 1. To determine whether the typicality of category exemplars affects infants' lexical processing. Specifically, is lexical processing more robust for referents that are typical category exemplars than referents that are atypical category exemplars? That is, are infants' word meanings initially narrow (faster and more accurate lexical processing for typical category exemplars) or broad (similar processing for typical and atypical category exemplars)?

Aim 2. To investigate how typicality effects change with age. A narrow-to-broad account would predict that the effect of category typicality on word recognition interacts with age, such that only more typical category referents are recognized reliably by younger infants, and the successful recognition of atypical category referents emerges with age. A broad-to-narrow account would predict early success at recognizing both typical and atypical category referents from an early age, with typicality differences emerging later in development as infants accrue additional experience with words and their possible referents (i.e., no interaction between typicality and age given the young age of infants in our current sample, or an interaction such that typicality differences emerge late).

Aim 3. To test whether individual differences between participants in word recognition and typicality effects are predicted by differences in experience (via parental report). Our third aim is to investigate whether individual differences in the effects of typicality are related to differences in infants' caregiver-reported experiences with the specific categories and exemplars tested in this experiment. A narrow-to-broad account of word meanings would predict that variability in children's early experience with category referents explains and therefore should predict the magnitude of typicality effects on word recognition, such that word recognition is more robust for an exemplar that is typical *for that child*. On a broad-to-narrow view, words are predicted to be extended to a broad class of category members even at a young age, such that word meanings should be less sensitive to

early experience with individual exemplars. This account thus predicts little to no overall effect of individual variation in early experience on word recognition for infants in our sample.¹

Methods

Participants

Families will be recruited from the United States through Lookit (Scott & Schulz, 2017; <https://lookit.mit.edu/>), an online infant testing platform. The study has received IRB approval to be conducted online using the Lookit platform. Parents sign up to be included in a participant database and are notified of their child's eligibility for studies via email. English-learning infants between 14 and 18 months of age ($N = 80$; 40 female infants planned; 410 days - 560 days, to include infants a few weeks younger than 14 months and a few weeks older than 18 months) that have no history of developmental concerns will participate in two sessions approximately one week apart. In pilot data using Lookit (see supplementary materials section S7), we found the following distribution of self-reported race/ ethnicity among our participants: approximately 60% White, 18% Asian, 3% Hispanic, 2% African American, and 17% reporting multiple categories. Families will receive gift cards for participating in the study.

The sample size was determined by conducting a power analysis using the jpower module in jamovi (The jamovi project, 2020) and considering resource constraints (Lakens, 2021). A sample size of $N=80$ provides 90% power to detect an effect size of $d=0.37$ (and 80% power to detect an effect size of $d=0.32$) in the main test of the typicality effect (Aim 1). The chosen sample size reflects the goal of striking a balance between having adequate power to detect effects large enough to be of theoretical interest and our data collection constraints for the project, based on the resources required for recruitment, compensation for participants, and video coding. We estimated the maximum number of

¹ As in Aim 2, a broad-to-narrow account predicts that experience gradually helps learners refine and constrain initially broad word meanings, such that individual differences in experience may play an increasing role in word recognition as children grow older. However, in Aim 3, our primary focus is on the overall effect of individual differences in exemplar experience on word recognition in our current sample (rather than its interaction with age).

participants we could recruit and test in the current experiment to be approximately 100 families. Assuming an estimated dropout rate of 20%, we arrived at the maximum sample size of $N=80$ and determined that this sample provides sufficient power to detect effect sizes within a range that is of most theoretical and practical significance (see Analysis Plan section below for an extended discussion of power with respect to the analyses of interest). A detailed list of exclusionary criteria at the participant level can be found in the supplementary materials (see S1.1. Participant-Level Exclusions). Infants that do not meet eligibility criteria will be replaced until a final sample of 80 participants is reached.

We chose the current age range (14-18 months) based on three main considerations. First, previous research suggests that there may be developmental changes in infants' word recognition for category exemplars varying in typicality and familiarity between 12 and 18 months of age (Garrison et al., 2020; Meints et al., 1999). Second, recent research suggests that there is a qualitative shift in infants' performance on looking-while-listening tasks at approximately 14 months (Bergelson, 2020). Given that we plan to evaluate individual differences in word recognition, we chose an age range during which infants' word recognition in this paradigm becomes more robust. Third, in our pilot data, we observed that infants below 14 months of age showed less robust word recognition than infants over 14 months of age, consistent with the quantitative shift observed in previous research (see S7.2 Online Pilot Experiment Results in the supplementary materials for further details).

Stimuli

All stimuli are available on the Open Science Framework (OSF) at https://osf.io/3t8gf/?view_only=d4031533ac8549d5a8682b7819ba7bee.

Animal Words/ Categories

We selected four basic-level animal categories for which to test infants' word recognition. The superordinate category of animals was selected because we expected that the items would elicit high engagement across this age range. The category labels will be a set of words (*dog, bird, cat, fish*)

selected according to the following criteria: (a) all words are understood at a relatively young age (see S2.1. Word Properties in the supplementary materials) and (b) exemplars from the corresponding basic-level category vary substantially in terms of typicality. All words occur frequently in corpora of child-directed speech in American English (see S2.1.).

Figure 1. Typical and atypical items for each category.

Images

The stimuli are a subset of 70 animal images collected from open-source image databases (Brodeur et al., 2010; Emberson & Rubinstein, 2016; Wikimedia Commons; <https://unsplash.com>). Images were edited to isolate the target animal and placed on a white 600 x 600 pixel background. For each category (*dog, cat, bird, fish*), we selected 16-20 images that varied in typicality. We then conducted a norming study with adults to collect objective ratings of typicality and nameability for all items within a category (see supplementary materials S2.2 and S3 for further details). The final set of items includes three typical and three atypical exemplars from each of the animal categories (n = 24) that were likely to be named by adults using the basic-level category label (see Figure 1; S2.2. in the

supplementary materials provides details on typicality ratings). Three images of naturalistic scenes (e.g., mountains) were selected to serve as attention getters.

Audio

A female native speaker of American English recorded 16 target sentences in child-directed speech. The sentences included four different carrier phrases: “Look at the [target label]”, “Find the [target label]”, “Do you see the [target label]”, “Where’s the [target label]”. Sentences were recorded in a sound-attenuated booth using Praat (Boersma & Weenink, 2019). All sentences were normalized to a total length of 1650 ms and an average intensity of 65 dB. 1500 ms of silence prior to the onset of the carrier phrase and 4000 ms of silence after the offset of the target word were appended to each recording, such that each auditory stimulus had a total duration of 7150 ms. Sentences were edited such that onset of the target label was consistent across all recordings (always 2650 ms after the onset of the auditory stimulus). Attention getters will be presented with 5200 ms of instrumental music.

Design & Procedure

Families will be alerted of their infants’ eligibility via the Lookit platform and will provide video documentation of consent. Before beginning the experiment, parents will be asked to complete an at-home set-up to maximize video recording quality by reducing backlighting and home distractions. Parents will seat their child comfortably in front of a laptop or desktop computer and close their eyes or turn around to limit their influence on infants’ looking behavior.

Infants will participate in two data collection sessions approximately one week apart. Word recognition for the four basic-level animal categories will be assessed using a looking-while-listening paradigm presented on Lookit (Scott & Schulz, 2017). During each session, infants will be presented with 24 trials with animal referents that vary in typicality. Attention getters will be distributed evenly throughout the experiment to maintain interest in the task.

Figure 2. Sample trial order for one block. Each target item is tested once per block with half of the trials displaying typical items and half displaying atypical items.

On each trial, infants are presented with two animal images that are matched on typicality and hear a phrase labeling one of the animals. For example, a typical trial for the target noun *bird* would present an image of a robin and an image of a tabby cat, while an atypical trial would present an image of a kingfisher and an image of a sheepdog (see Figure 2). Pairs of items from different basic-level categories will be yoked together (e.g., sparrow will always appear with beagle) such that each basic-level category will occur equally often with each of the other basic-level categories. Each item will appear twice as the target and twice as the distractor for a total of four exposures across two sessions.

Trials will be presented in six blocks of four separated by an attention getter, for a total of 24 trials per session. Within a block, infants will see two typical and two atypical trials testing each of the four target labels (*bird*, *cat*, *dog*, *fish*) displayed in a pseudorandom order so the same target item does not appear more than twice in a row (Figure 2). The order of trial presentation and the target location will be counterbalanced across participants and testing sessions. Within a single session, each basic-level category label will be tested six times and the location of the target item will appear equally often

on the left and right side of the screen. For the second session, both the order of trial presentation and target location of the category label will be reversed (see the OSF page for sample trial orders).

Parental Survey. Parents will complete a survey after their infant completes the second testing session. The survey will be composed of two components - a measure of children's vocabulary size (MB-CDI Level I Short Form, slightly modified to include all four category labels tested in the study; Fenson et al., 2000) and a survey on infants' prior experience with each category exemplar presented during the experiment. Parents will judge each of the 24 items presented during the experiment in terms of typicality for their child (*experience-based typicality*). For each image, parents will be asked "How typical is this image of the [category label] your child experiences?" and rate the image on a scale from 1 (very atypical) to 5 (very typical). Parents will also be asked to report the contexts in which their infant experiences each animal (e.g., as a pet at home, in a book you read, at the local zoo, etc.). The full instructions for the survey can be found in the supplementary materials (see S4).

Video Coding. Infants' eye-gaze behavior will be coded frame-by-frame from webcam video recordings. Details on our procedure for coding participant videos and assessing coding reliability can be found in the supplementary materials (S5).

Pilot Data. We conducted a pilot experiment to test the feasibility of our design. In the supplementary materials, we include details on the experiment methods (S7.1) and preliminary results demonstrating that we can successfully measure word recognition on Lookit (S7.2).

Analysis Plan

Overview of Analytic Approach

Modeling framework. All data processing and analysis will be conducted using software packages in R (R Development Core Team, 2021). Linear mixed-effects models will be fit using the lme4 package (Bates et al., 2015); p -values for mixed-effects models will be computed using the lmerTest package (Kuznetsova et al., 2017), which uses the Satterthwaite approximation to estimate degrees of freedom for statistical tests. Predictor variables will be centered unless otherwise specified.

For each mixed-effects model specified below, we will first fit a model with the maximal random-effects structure at the participant level (Barr et al., 2013). In models fit to trial-level looking data, we will also include random intercepts for target word to account for non-independence due to items. In cases where the models with the specified random-effects structure fail to converge, we will iteratively prune random effects until convergence is achieved, removing random effects of lesser theoretical interest prior to removing random effects of greater theoretical interest. Specifically, we will first remove random effects for covariates, then random effects for covariances among random effects (beginning with covariances close to zero), followed by random intercepts for target word and then for participant (Brauer & Curtin, 2018). Random slopes for the effects of interest (typicality condition and experience-based typicality) will be removed only if all other pruning efforts have been unsuccessful.

Windows of analysis. In our main analyses, we will investigate the difference between children's proportion looking to the target image during a *critical window* after the onset of the target word and a *baseline window* prior to the onset of the target word. The baseline window will be set to encompass the 2s immediately preceding the target word onset (-2000 ms – 0 ms), consistent with past literature (e.g., Garrison et al., 2020). The size of the critical window varies across the word recognition literature, with some experiments analyzing shorter windows of 300 ms – 1800 ms or 367 ms – 2000 ms post target word onset (Fernald et al., 2008; Garrison et al., 2020; Swingley & Aslin, 2002) and other experiments analyzing longer time windows extending as long as 3500 ms post target word onset (Bergelson & Swingley, 2012; Meints et al., 1999). The choice in window selection carries a tradeoff between limiting the focus of the analysis on infant behavior that is most closely related to the onset of the target word (in the case of the smaller windows of analysis), on the one hand, and capturing as much relevant looking behavior as possible, on the other. In particular, recent research suggests that longer time windows may increase reliability, a key concern in estimating individual differences (Zettersten et al., 2021). Given that we will test a wide age range in the current study, we will set our critical window to lie between 300ms-2800ms, as a compromise between these tradeoffs.

We will also investigate the degree to which any results depend on the choice of a longer critical time window (see robustness analysis section below).

Trial-level data processing and exclusion. Only trials with sufficient looking data (>50% looking during both the critical window and the baseline window) and without procedural errors will be included in analyses. Trial-level exclusion criteria can be found in the supplementary materials (S1.2. Trial-Level Exclusions).

Dependent measure: Baseline-corrected proportion target looking. We will compute proportion of looking at the target item as opposed to the distractor item for both the baseline window (-2000ms - 0ms) and the critical window (300ms - 2800ms): $\text{target}/(\text{target}+\text{distractor})$. We will then compute a baseline-corrected proportion target looking measure, by subtracting proportion looking to the target during the baseline window from proportion looking to the target during the critical window for each trial. Baseline correction is used to account for potential saliency differences between target and distractor images (e.g., Garrison et al., 2020; Meints et al., 1999). For models testing the average typicality effect at the participant level (Aim 1, section 1.1), we will compute the average baseline-corrected proportion target looking for both typical and atypical items. For models fit to trial-level data (i.e., the model in section 1.2 and the models in Aims 2 and 3), models will be fit to baseline-corrected proportion target looking for each individual trial.

Aim 1: To determine whether the typicality of category exemplars affects infants' lexical processing.

1.1. Average participant-level analysis

To investigate whether there is an effect of typicality on infants' word recognition, we will first conduct an overall analysis averaging across all items and categories that participants encounter during the study. For each infant, we will compute the average baseline-corrected proportion of looking to the target within the critical window across trials for each typicality condition (i.e., each infant will provide two baseline-corrected proportion looking scores: an average baseline-corrected target looking

score for typical items and an average baseline-corrected target looking score for atypical items). We will report the mean and 95% CI of baseline-corrected target looking for each typicality condition as descriptive statistics. To test our questions of interest, we will fit a linear mixed-effects model predicting average baseline-corrected proportion target looking from typicality condition (centered), including a by-participant random intercept and random slope for typicality condition. The model will be specified as follows:

$$\text{average_corrected_target_looking} \sim 1 + \text{typicality_condition} + (1 + \text{typicality_condition} | \text{participant})$$

Using this model, we will answer the following three questions:

1.1.1. Do infants successfully recognize the target word? (quality check)

Chance-level baseline-corrected proportion target looking is 0%. The intercept of the model will therefore test whether infants successfully recognize the target words (averaging across typicality conditions, because typicality condition will be centered). Testing whether infants show successful word recognition overall will serve as a quality check of our method, ensuring that infants are engaging in the task as expected. We will treat infants successfully recognizing the target words on average as a basis for interpreting effects of typicality: typicality effects will only be interpreted as meaningful if infants show above-chance word recognition collapsing across typicality conditions. Our sample size ($N=80$) ensures that we will have over 99% power to detect an effect of $d=0.5$ or larger. An effect size of $d=0.5$ is smaller than the effect typically found in studies of online familiar word recognition (meta-analytic effect size: $d=1.24$; Bergmann et al., 2018), including for infants between 14-18 months of age (e.g., Garrison et al., 2020: $d \sim 0.5 - 1$ across an age range of 12-18 months).

1.1.2. Do infants differentially recognize typical vs. atypical items?

The main effect of typicality is statistically equivalent to a paired-samples t -test and will test whether there are differences in infants' recognition of typical items compared to atypical items. Given our sample size of $N=80$ participants, we will have 90% power to detect an effect size of $d=0.37$ (80% power to detect an effect size of $d=0.32$). Past studies most similar to the current design have reported

a wide spectrum of absolute effect sizes ranging from $d = 0.03$ to $d = 0.64$ (see Supplementary Materials S6 for an overview; Garrison et al., 2020; Meints et al., 1999; Perry & Saffran, 2017; Southgate & Meints, 2000). Our design is well-powered to detect an effect consistent with the upper range of effect sizes observed in past studies, but has limited power to detect small effect sizes in the lower range (see below for our rationale in choosing the smallest effect size of interest). A significant effect in this analysis will lead us to conclude that the typicality of category exemplars makes a difference in how easily infants recognize familiar words. Specifically, we predict that infants will show a higher average baseline-corrected proportion looking to the target item when it is typical than when it is atypical.

A non-significant effect in this analysis will not allow us to infer that there is no effect of typicality on infants' looking behavior. Instead, we will use equivalence testing to assess whether the effect is smaller than a meaningful effect size threshold, sometimes called the smallest effect size of interest (Lakens et al., 2018). We selected our smallest effect size of interest using the small-telescopes approach (Simonsohn, 2015), according to which the lower bound for equivalence testing is set as the effect size earlier studies would have had 33% power to detect. The most relevant studies from past literature that we identified have used sample sizes of approximately 30-40 participants (Garrison et al., 2020; Meints et al., 1999; Perry & Saffran, 2017). The smallest effect size that studies with this sample size would have 33% power to detect is approximately $d=0.25$, which we choose as our lower threshold for an effect size of interest. If the typicality effect is not significant, we will conduct an equivalence test evaluating whether the null hypothesis that the absolute effect size is at least as large as $d=0.25$ can be rejected, and the alternative hypothesis that the effect size lies between -0.25 and 0.25 is accepted. In other words, this test will allow us to conclude whether a non-significant effect is sufficiently small to be considered practically equivalent to a null effect.

1.1.3. Do infants successfully recognize words in the typical condition and in the atypical condition?

We will further test whether infants successfully recognized words in each typicality condition within the same model, by recoding the typicality condition predictor in the model to be centered on the typical condition (atypical = -1; typical=0) and on the atypical condition (atypical = 0; typical = 1). For each of these re-centered models, the intercept will now indicate whether infants' baseline-corrected looking during the critical window was significantly above chance for each of the typicality conditions. This design provides us with at least 80% power to detect an effect size of $d=0.32$ or larger. As in the analysis above, if either effect is non-significant, we will conduct an equivalence test with a threshold of $d=0.25$ to assess whether the effect size is statistically equivalent to an effect of minimal interest.

1.2. Trial-level analysis

To further test the effect of typicality, we will also fit a linear mixed-effects model predicting trial-by-trial baseline-corrected proportion target looking from typicality condition (centered), including a by-participant random intercept and random slope for typicality condition as well as a random intercept for target word (i.e., *dog*, *cat*, etc.). Modeling trial-level looking data while accounting for non-independence due to the target word allows us to assess the degree to which typicality effects generalize across particular category labels. The model will be specified as follows:

$$\text{corrected_target_looking} \sim 1 + \text{typicality_condition} + \\ (1 + \text{typicality_condition} | \text{participant}) + \\ (1 | \text{target_word})$$

To estimate our statistical power, we used the *simr* package in R (Green & Macleod, 2016) to compute power for a range of possible effect sizes. Specifically, we fit the model above using the trial-by-trial looking data from our pilot study, and subsequently simulated model outcomes for a range of fixed effect sizes for typicality, assuming a sample size of $N=80$ and similar variability in looking behavior as in the pilot. These estimates likely represent an underestimate of the power in our current experimental design, which includes three times the number of trials (48 total trials rather than 16 in

the pilot experiment). We estimate that our power ranges between (at least) 79%-99% for an effect of typicality ranging from $b = 0.06$ to $b = 0.10$. This effect size range would reflect a 6%-10% difference in the typicality effect (difference in baseline-corrected proportion target looking between typical and atypical items).

In general, we expect to observe consistent results between the average participant-level analysis in model 1.1. and the trial-level analysis in model 1.2. Any discrepancies will be interpreted as indicating that results depend on whether item-level non-independence is accounted for in the modeling approach (the main conceptual difference between the two models). For example, if a significant typicality effect is found in the participant-level model from section 1.1, but not in the trial-level model from section 1.2, this will be interpreted as indicating that the typicality effect does not generalize across items (i.e., category labels). In a scenario such as this, we plan to conduct exploratory analyses to investigate the source of the difference in outcomes when by-item random effects are included in the analytic approach (e.g., inspecting item-level variation in target looking).

1.3. Timecourse analysis

While the previous analyses test overall differences in infants' accuracy in recognizing typical and atypical category exemplars, they offer limited insight into how infants' processing unfolds across time. In order to gain finer-grained insight into the timecourse of infants' word recognition for typical compared to atypical items, we will conduct a cluster-based permutation analysis (Maris & Oostenveld, 2007). This analysis will identify where over the course of a test trial looking to the target diverges between the two typicality conditions, while controlling Type I error rate. The cluster-based permutation analysis will be conducted using the *eyetrackingR* package (Dink & Ferguson, 2015). The looking data will first be grouped into 67ms time bins across a wider time window ranging from -2000 ms to 4000 ms relative to target word onset (since infants in the pilot study generally tend to begin disengaging from the screen at approximately 4000ms after the target word onset; 67ms corresponds to the lowest trial-level sampling rate that we include in our analysis, 15 Hz). Since this analysis allows

us to assess differences in looking between typical and atypical items across the entire timecourse, including during the baseline window, we analyze participants' average proportion target looking at each time bin in the window of interest (i.e., without introducing any correction for baseline looking). *T*-tests comparing performance between conditions are computed within each time bin, and adjacent time bins with a significant *t*-statistic (at an alpha level of 0.05) will be grouped into a cluster. Next, the *t*-statistics for a given cluster will be summed to create a statistic representing the overall cluster size in which a difference was observed. Then, we will test the probability of observing a cluster of the same size in a bootstrapping analysis with 1,000 simulations. Within each simulation, the condition labels (typical vs. atypical) will be randomly shuffled and clusters will be computed using the same method as above. The simulations will yield a distribution of cluster sizes based on chance assignment of condition labels, allowing us to evaluate the probability of obtaining clusters indexing condition differences of the same size as those observed in our data. While power is difficult to estimate *a priori* for this method given uncertainty surrounding how target looking may differ across the timecourse, our sample size is at least as large as many infant studies that have used this method in the past (e.g., Kartushina & Mayor, 2019; LaTourrette & Waxman, 2019).

Aim 2. To investigate how typicality effects change with age.

To investigate whether the overall effect of typicality depends on participant age, we will extend the trial-level linear mixed-effects model from 1.2. to include an interaction with participant age. We will fit a linear mixed-effects model predicting participants' trial-by-trial baseline-corrected proportion target looking from typicality condition, participant age (mean-centered) and their interaction, including by-participant and by-target word random intercepts and a by-participant random slope for typicality condition. A significant interaction between typicality condition and age such that the typicality effect decreases with age would provide evidence consistent with a narrow-to-broad account of early word representations, suggesting that word representations are initially focused on a narrow set of category exemplars and increase in generality with age. Non-significant interactions (or

interactions in the opposite direction, e.g., increasing effects of typicality with age) will be interpreted as evidence against a narrow-to-broad account of early word representations. The model will be specified as follows:

$$\text{corrected_target_looking} \sim 1 + \text{typicality_condition} * \text{age} + \\ (1 + \text{typicality_condition} | \text{participant}) + \\ (1 | \text{target_word})$$

To estimate our statistical power, we used the same simulation-based approach as in section 1.2. We fit the interaction model above using the trial-level looking data from our pilot study and subsequently simulated model outcomes for a range of possible interaction effect sizes, assuming a sample size of $N=80$ and similar variability in looking behavior as in the pilot. We estimate that our power ranges between (at least) 32%-86% for an interaction effect size ranging from $b = 0.025$ to $b = 0.05$. This effect size range would reflect a 2.5%-5% increase in the typicality effect (difference in baseline-corrected proportion target looking between typical and atypical items) for each month of age.

Aim 3. To test whether individual differences between participants in word recognition and typicality effects are predicted by differences in experience.

To address the third aim, we will investigate whether individual differences in children's word recognition is related to caregiver report of infants' typical experience with specific category exemplars, while controlling for age. In this analysis, the focus is on testing whether experience-based typicality is related to the typicality effect on participant accuracy. We will fit a linear mixed-effects model predicting participants' trial-by-trial baseline-corrected proportion target looking from experience-based typicality (for a given target image/ category exemplar) and age (mean-centered). We will include random intercepts for participant and target word as well as a by-participant random slope for experience-based typicality. The item-level experience-based typicality measure will be z -scored within a given participant. The model will be specified as follows:

$$\text{corrected_target_looking} \sim 1 + \text{item_experience_typicality} + \text{age} + \\ (1 + \text{item_experience_typicality} | \text{participant}) + \\ (1 | \text{target_word})$$

Power is difficult to estimate for this analysis due to uncertainty in variance in caregivers' experience-based typicality ratings for their infants (not assessed in our pilot study). As a proxy, we estimated power in this design based on the participant-level variability we observed in our typicality norming study, randomly simulating experience-based typicality responses for each participant using the distributions of typicality norms for each image. While we expect the caregivers' typicality ratings to differ from adult typicality ratings, the distribution of experience-based typicality ratings is likely to be comparable to distributions for adult typicality ratings, such that the typicality rating distributions are a suitable proxy for power simulations. We then used the *simr* package in R (Green & Macleod, 2016) to estimate our power to detect a range of typicality effect sizes in the full model specified below (including age as a covariate and the specified random effects structure), by simulating model outcomes at our target sample size based on the variance in looking behavior observed in our pilot data and simulated typicality ratings (5,000 simulations for each model and simulated effect size). These estimates likely represent an underestimate of the proposed design's power, given that the proposed design includes significantly more trials (48 trials as compared to 16 trials per participants) compared to the pilot study that we use as a basis for the simulations (see S7). Based on these simulations, we estimated that our power ranges between 54%-99% for effect sizes between $b = 0.025$ and $b = 0.05$, which corresponds to predicting a 5-10% difference in proportion target looking when moving from -1 SD (atypical) to +1 SD (typical) along the experience-based typicality predictor. In other words, if experience-based typicality exerts a large effect on infants' word processing, we are well-powered to detect such an effect; however, if the influence of experience is small, the current design is underpowered, meaning that we will interpret non-significant effects with caution.

4. Robustness analyses

We will also conduct a series of robustness analyses to probe the degree to which any results hinge on key analytic decisions. In interpreting any discrepancies in results between alternative

analyses, we will focus our interpretation on the primary analyses, but qualify results if they are not robust across the analytic dimensions outlined below.

4.1. Window choice. We will conduct the same analyses as those described above using an alternative, shorter critical window of 300-1800ms (e.g., Fernald et al., 2008) (with the exception of the cluster-based permutation analysis described in analysis section 1.3).

4.2. Excluding trials with potentially unknown words. We will also repeat the analyses above after excluding all trials with target labels that infants are not marked as comprehending in corresponding caregivers' MCDI responses. This analysis is to ensure that any absence of typicality differences cannot be attributed to infants in our sample simply not comprehending the (putatively) familiar words.

4.3. Alternative dependent measures in Aim 1: reaction time. To assess the consistency of typicality effects across multiple indices of word recognition, we will fit the main two models in Aim 1 (section 1.1 and 1.2 in the Analysis Plan), testing for a typicality effect at the participant level and at the trial level, using a second dependent measure: reaction time. In looking-while-listening studies, reaction time is typically defined as the time it takes for infants to first shift their gaze to the target image within the critical window. In this measure, we will follow the approach of Fernald et al. (2008) and specifically select all trials in which an infant is looking at the distractor at the onset of the target label. Then, for each trial, we will compute the amount of time from the onset of the target label until infants' first look at the target object within the critical window of 300 ms - 2800 ms (given that they were initially looking at the distractor item). Since reaction times typically exhibit positive skew, we will log-transform trial-level reaction times (following Donnelly & Kidd, 2020) and compute by-participant average log reaction times for each typicality condition. Only infants who contribute at least 4 typical and 4 atypical distractor-to-target reaction time trials (after applying trial-level exclusions) will be included in the reaction time analyses.

4.4. Testing whether effects interact with test session. Given the multi-session structure of our data collection procedure, we will also test whether analyses hold individually in each test session, by

repeating the analyses above while including an interaction with test session (session 1 vs. session 2). We do not predict there to be any strong differences between sessions, either in overall accuracy or in the magnitude of the typicality effect; however, this analysis will serve to ensure that there are no systematic differences depending on infants' experience with the test paradigm, as has been observed in some infant studies (e.g., Santolin et al., 2021). If such differences are observed, we will place a greater emphasis on results observed in the first test session, at which point infants would not yet have become familiar with the test or stimulus structure.

Timeline

The first submission of the Stage 1 manuscript occurred in mid-January 2021 and the second resubmission will occur at the end of December 2021. Assuming the review and revision process takes an additional 2-3 months in total (in-principle acceptance February/ March 2022), we propose the following (conservatively estimated) timeline for data collection and resubmission of the Stage 2 manuscript, pending in-principle acceptance. Should the review and revision process extend beyond February 2022, all estimates would be shifted relative to the final date of in-principle acceptance.

February/ March 2022 – May/ June 2022: Data collection phase

During piloting, we averaged approximately 15 participants per week collected through the Lookit platform. Adjusting for potential exclusions and additional delays introduced by testing participants in two sessions, the current plan allows for up to 16 weeks of data collection, which would correspond to a rate of 11 participants per week assuming an exclusion rate of 30% of sessions. If needed, we could increase the speed of data collection by independently recruiting families and directing them to the Lookit platform to participate in the study.

June 2022 - August 2022: Final coding of participant videos

Hand-coding of participant videos will require substantial resources. This timeline assumes that participant videos will be coded continuously as data collection unfolds but allows for additional time for coding after data collection has completed.

August 2022 - October 2022: Analysis and finalizing Stage 2 manuscript for submission

We allow 2-3 months for conducting the planned analyses, discussing and conducting potential post-hoc, exploratory analyses, and writing the discussion.

References

- Ambridge, B. (2020). Against stored abstractions: A radical exemplar model of language acquisition. *First Language*, 40(5–6), 509–559. <https://doi.org/10.1177/0142723719869731>
- Arias-Trejo, N., & Plunkett, K. (2010). The effects of perceptual similarity and category membership on early word-referent identification. *Journal of Experimental Child Psychology*, 105(1–2), 63–80. <https://doi.org/10.1016/j.jecp.2009.10.002>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Barrett, M. D. (1986). Early semantic representations and early word-usage. In S. A. Kuczaj & M. D. Barrett (Eds.), *The Development of Word Meaning* (pp. 39–67). Springer.
- Bates, D., Mächler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Bauer, P. J., Dow, G. A., & Hertsgaard, L. A. (1995). Effects of prototypicality on categorization in 1- to 2-year-olds: Getting down to basic. *Cognitive Development*, 10(1), 43–68. [https://doi.org/10.1016/0885-2014\(95\)90018-7](https://doi.org/10.1016/0885-2014(95)90018-7)
- Beckage, N., Smith, L. B., & Hills, T. (2011). Small worlds and semantic network growth in typical and late talkers. *PLoS ONE*, 6(5). <https://doi.org/10.1371/journal.pone.0019348>
- Bergelson, E. (2020). The comprehension boost in early word learning: Older infants are better learners. *Child Development Perspectives*, 14(3), 142–149. <https://doi.org/10.1111/cdep.12373>
- Bergelson, E., & Aslin, R. (2017). Semantic specificity in one-year-olds' word comprehension. *Language Learning and Development*, 13(4), 481–501. <https://doi.org/10.1080/15475441.2017.1324308>
- Bergelson, E., & Swingley, D. (2012). At 6-9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences of the United States of*

America, 109(9), 3253–3258. <https://doi.org/10.1073/pnas.1113380109>

Bergmann, C., Tsuji, S., Piccinini, P. E., Lewis, M. L., Braginsky, M., Frank, M. C., & Cristia, A. (2018). Promoting replicability in developmental research through meta-analyses: Insights from language acquisition research. *Child Development*, 89(6), 1996–2009.

<https://doi.org/10.1111/cdev.13079>

Bleses, D., Makransky, G., Dale, P. S., Højen, A., & Ari, B. A. (2016). Early productive vocabulary predicts academic achievement 10 years later. *Applied Psycholinguistics*, 37(6), 1461–1476.

<https://doi.org/10.1017/S0142716416000060>

Boersma, P., & Weenink, D. (2019). *Praat: doing phonetics by computer*. <http://www.praat.org>

Bornstein, M. H., & Arterberry, M. E. (2010). The development of object categorization in young children: Hierarchical inclusiveness, age, perceptual attribute, and group versus individual analyses. *Developmental Psychology*, 46(2), 350–365. <https://doi.org/10.1037/a0018411>

Borovsky, A., Ellis, E. M., Evans, J. L., & Elman, J. L. (2016). Lexical leverage: Category knowledge boosts real-time novel word recognition in 2-year-olds. *Developmental Science*, 19(6), 918–932.

<https://doi.org/10.1111/desc.12343>

Borovsky, A., & Peters, R. E. (2019). Vocabulary size and structure affects realtime lexical recognition in 18-month-olds. *PLoS ONE*, 14(7), 1–21. <https://doi.org/10.1371/journal.pone.0219290>

Braginsky, M., Yurovsky, D., Marchman, V. A., & Frank, M. C. (2019). Consistency and variability in children's word learning across languages. *Open Mind*, 3, 52–67.

https://doi.org/10.1162/opmi_a_00026

Brauer, M., & Curtin, J. J. (2018). Linear mixed-effects models and the analysis of nonindependent data: A unified framework to analyze categorical and continuous independent variables that vary within-subjects and/or within-items. *Psychological Methods*, 23(3), 389–411.

<https://doi.org/10.1037/met0000159>

Brodeur, M. B., Dionne-Dostie, E., Montreuil, T., & Lepage, M. (2010). The bank of standardized

- stimuli (BOSS), a new set of 480 normative photos of objects to be used as visual stimuli in cognitive research. *PLoS ONE*, 5(5), e10773. <https://doi.org/10.1371/journal.pone.0010773>
- Colunga, E., & Sims, C. E. (2017). Not only size matters: Early-talker and late-talker vocabularies support different word-learning biases in babies and networks. *Cognitive Science*, 41, 73–95. <https://doi.org/10.1111/cogs.12409>
- De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2019). The “Small World of Words” English word association norms for over 12,000 cue words. *Behavior Research Methods*, 51(3), 987–1006. <https://doi.org/10.3758/s13428-018-1115-7>
- Dink, J. W., & Ferguson, B. (2015). *eyetrackingR*. <http://www.eyetracking-r.com>
- Donnelly, S., & Kidd, E. (2020). Individual differences in lexical processing efficiency and vocabulary in toddlers: A longitudinal investigation. *Journal of Experimental Child Psychology*, 192, 104781. <https://doi.org/10.1016/j.jecp.2019.104781>
- Eimas, P. D., & Quinn, P. C. (1994). Studies on the formation of perceptually based basic-level categories in young infants. *Child Development*, 65(3), 903–917. <https://doi.org/10.2307/1131427>
- Emberson, L. L., & Rubinstein, D. Y. (2016). Statistical learning is constrained to less abstract patterns in complex sensory input (but not the least). *Cognition*, 153, 63–78. <https://doi.org/10.1016/j.cognition.2016.04.010>
- Fenson, L., Pethick, S., Renda, C., Cox, J. L., Dale, P. S., & Reznick, J. S. (2000). Short-form versions of the MacArthur Communicative Development Inventories. *Applied Psycholinguistics*, 21, 95–115. <https://doi.org/10.1017/S0142716400001053>
- Ferguson, B., & Waxman, S. (2017). Linking language & categorization in infancy. *Journal of Child Language*, 44(3), 527–552. <https://doi.org/10.1017/S0305000916000568>
- Fernald, A., Zangl, R., Portillo, A. L., & Marchman, V. A. (2008). Looking while listening: Using eye movements to monitor spoken language comprehension by infants and young children. In I. A. Sekerina, E. M. Fernandez, & H. Clahsen (Eds.), *Developmental Psycholinguistics: On-line*

Methods in Children's Language Processing (pp. 97–135). John Benjamins.

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

Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2021). *Variability and consistency in early language learning: The Wordbank project*. MIT Press.

Garrison, H., Baudet, G., Breitfeld, E., Aberman, A., & Bergelson, E. (2020). Familiarity plays a small role in noun comprehension at 12–18 months. *Infancy*, 25(4), 458–477. <https://doi.org/10.1111/infa.12333>

Goldberg, A. (2019). *Explain me this: Creativity, competition, and the partial productivity of constructions*. Princeton University Press.

Green, P., & Macleod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493–498. <https://doi.org/10.1111/2041-210X.12504>

Gruendel, J. M. (1977). Referential extension in early language development. *Child Development*, 48(4), 1567–1576.

Hills, T., Maouene, M., Maouene, J., Sheya, A., & Smith, L. B. (2009). Longitudinal analysis of early semantic networks. *Psychological Science*, 20(6), 729–739. <https://doi.org/10.1111/j.1467-9280.2009.02365.x>

Hirsh-Pasek, K., Golinkoff, R. M., Hennon, E. A., & McGuire, M. J. (2004). Hybrid theories at the frontier of developmental psychology: The emergentist coalition model of word learning as a case in point. In D. G. Hall & S. R. Waxman (Eds.), *Weaving a lexicon* (pp. 173 – 204). MIT Press.

Hurley, K., & Oakes, L. M. (2018). Infants' daily experience with pets and their scanning of animal faces. *Frontiers in Veterinary Science*, 5, 1–11. <https://doi.org/10.3389/fvets.2018.00152>

Huttenlocher, J., & Smiley, P. (1987). Early word meanings: The case of object names. *Cognitive*

Psychology, 19(1), 63–89. [https://doi.org/10.1016/0010-0285\(87\)90004-1](https://doi.org/10.1016/0010-0285(87)90004-1)

Kartushina, N., & Mayor, J. (2019). Word knowledge in six- to nine-month-old Norwegian infants? Not without additional frequency cues. *Royal Society Open Science*, 6, 180711. <https://doi.org/10.1098/rsos.180711>

Kovack-Lesh, K. A., McMurray, B., & Oakes, L. M. (2014). Four-month-old infants' visual investigation of cats and dogs: Relations with pet experience and attentional strategy. *Developmental Psychology*, 50(2), 402–413. <https://doi.org/10.1037/a0033195>

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13). <https://doi.org/10.18637/jss.v082.i13>

Lakens, D. (2021). Sample size justification. *PsyArXiv*, 1–31. <https://doi.org/10.31234/osf.io/9d3yf>

Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science*, 1(2), 259–269. <https://doi.org/10.1177/2515245918770963>

LaTourrette, A., & Waxman, S. R. (2019). A little labeling goes a long way: Semi-supervised learning in infancy. *Developmental Science*, 22(1), 1–9. <https://doi.org/10.1111/desc.12736>

Mani, N., Johnson, E., McQueen, J. M., & Huettig, F. (2013). How yellow is your banana? Toddlers' language-mediated visual search in referent-present tasks. *Developmental Psychology*, 49(6), 1036–1044. <https://doi.org/10.1037/a0029382>

Marchman, V. A., Loi, E. C., Adams, K. A., Ashland, M., Fernald, A., & Feldman, H. M. (2018). Speed of language comprehension at 18 months old predicts school-relevant outcomes at 54 months old in children born preterm. *Journal of Developmental & Behavioral Pediatrics*, 39(3), 246–253. <https://doi.org/10.1097/DBP.0000000000000541>

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190. <https://doi.org/10.1016/j.jneumeth.2007.03.024>

- Mayor, J., & Plunkett, K. (2010). A neurocomputational account of taxonomic responding and fast mapping in early word learning. *Psychological Review*, 117(1), 1–31.
<https://doi.org/10.1037/a0018130>
- Meints, K., Plunkett, K., & Harris, P. L. (1999). When does an ostrich become a bird? The role of typicality in early word comprehension. *Developmental Psychology*, 35(4), 1072–1078.
<https://doi.org/10.1037/0012-1649.35.4.1072>
- Oakes, L. M., Coppage, D. J., & Dingel, A. (1997). By land or by sea: the role of perceptual similarity in infants' categorization of animals. *Developmental Psychology*, 33(3), 396–407.
<https://doi.org/10.1037/0012-1649.33.3.396>
- Perry, L. K., & Saffran, J. R. (2017). Is a pink cow still a cow? Individual differences in toddlers' vocabulary knowledge and lexical representations. *Cognitive Science*, 41(4), 1090–1105.
<https://doi.org/10.1111/cogs.12370>
- Perry, L. K., & Samuelson, L. K. (2011). The shape of the vocabulary predicts the shape of the bias. *Frontiers in Psychology*, 2(345), 1–12. <https://doi.org/10.3389/fpsyg.2011.00345>
- Peters, R., & Borovsky, A. (2019). Modeling early lexico-semantic network development: Perceptual features matter most. *Journal of Experimental Psychology: General*, 148(4), 763–782.
<https://doi.org/10.1037/xge0000596.supp>
- Poulin-Dubois, D., & Sissons, M. E. (2002). Is this still called a dog? 18-month-olds' generalization of familiar labels to unusual objects. *Infant and Child Development*, 11(1), 57–67.
<https://doi.org/10.1002/icd.256>
- Quinn, P. C., & Tanaka, J. W. (2007). Early development of perceptual expertise: Within-basic-level categorization experience facilitates the formation of subordinate-level category representations in 6- to 7-month-old infants. *Memory and Cognition*, 35(6), 1422–1431.
<https://doi.org/10.3758/BF03193612>
- R Development Core Team. (2021). *R: A language and environment for statistical computing*. R

Foundation for Statistical Computing. <https://www.r-project.org/>

Rescorla, L. A. (1980). Overextension in early language development. *Journal of Child Language*, 7(2), 321–335. <https://doi.org/10.1017/S0305000900002658>

Robinson, C. W. (2002). *Comprehension of labels denoting typical and atypical exemplars: What is guiding lexical extensions?* The University of Toledo.

Rosch, E. (1978). Principles of categorization. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings* (pp. 189–206). MIT Press.

Santolin, C., Garcia-Castro, G., Zettersten, M., Sebastian-Galles, N., & Saffran, J. R. (2021).

Experience with research paradigms relates to infants' direction of preference. *Infancy*, 26(1), 39–46. <https://doi.org/10.1111/infa.12372>

Scott, K., & Schulz, L. (2017). Lookit (Part 1): A new online platform for developmental research. *Open Mind*, 1(1), 4–14. https://doi.org/10.1162/opmi_a_00002

Simonsohn, U. (2015). Small telescopes: Detectability and the evaluation of replication results. *Psychological Science*, 26(5), 559–569. <https://doi.org/10.1177/0956797614567341>

Smith, L. B., Jones, S. S., Landau, B., Gershkoff-Stowe, L., & Samuelson, L. (2002). Object name learning provides on-the-job training for attention. *Psychological Science*, 13(1), 13–19.

Snedeker, J., & Gleitman, L. (2004). Why it is hard to label our concepts. In D. G. Hall & S. R. Waxman (Eds.), *Weaving a Lexicon* (pp. 257–294). MIT Press.

Southgate, V., & Meints, K. (2000). Typicality, naming, and category membership in young children. *Cognitive Linguistics*, 11(1/2), 5–16. <https://doi.org/10.1515/cogl.2001.011>

Steyvers, M., & Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive Science*, 29(1), 41–78.
https://doi.org/10.1207/s15516709cog2901_3

Swingley, D., & Aslin, R. N. (2002). Lexical neighborhoods and the word-form representations of 14-month-olds. *Psychological Science*, 13(5), 480–484. <https://doi.org/10.1111/1467-9280.00485>

The jamovi project. (2020). *jamovi* (1.2). <https://www.jamovi.org>

Tincoff, R., & Jusczyk, P. W. (1999). Some beginnings of word comprehension in 6-month-olds.

Psychological Science, 10(2), 172–175. <https://doi.org/10.1111/1467-9280.00127>

Waxman, S. R., & Gelman, S. A. (2009). Early word-learning entails reference, not merely

associations. *Trends in Cognitive Sciences*, 13(6), 258–263.

<https://doi.org/10.1016/j.tics.2009.03.006>

Waxman, S. R., & Markow, D. B. (1995). Words as invitations to form categories: Evidence from 12- to 13-month-old infants. *Cognitive Psychology*, 29(3), 257–302.

<https://doi.org/10.1006/cogp.1995.1016>

Willits, J. A., Wojcik, E. H., Seidenberg, M. S., & Saffran, J. R. (2013). Toddlers activate lexical semantic knowledge in the absence of visual referents: Evidence from auditory priming. *Infancy*,

18(6), 1–23. <https://doi.org/10.1111/infa.12026>

Wojcik, E. H. (2018). The development of lexical–semantic networks in infants and toddlers. *Child*

Development Perspectives, 12(1), 34–38. <https://doi.org/10.1111/cdep.12252>

Wojcik, E. H., & Saffran, J. R. (2013). The ontogeny of lexical networks: Toddlers encode the relationships among referents when learning novel words. *Psychological Science*, 24(10), 1898–

1905. <https://doi.org/10.1177/0956797613478198>

Yin, J., & Csibra, G. (2015). Concept-based word learning in human infants. *Psychological Science*,

26(8), 1316–1324. <https://doi.org/10.1177/0956797615588753>

Zettersten, M., Bergey, C., Bhatt, N., Boyce, V., Braginsky, M., Carstensen, A., DeMayo, B.,

Kachergis, G., Lewis, M., Long, B., MacDonald, K., Mankewitz, J., Meylan, S. C., Saleh, A. N.,

Schneider, R. M., Tsui, A., Uner, S., Xu, T. L., Yurovsky, D., & Frank, M. C. (2021). Peekbank:

Exploring children’s word recognition through an open, large-scale repository for developmental

eye-tracking data. *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*.

Supplementary Materials

S1. Participant-Level and Trial-Level Exclusions

S1.1. Participant-Level Exclusions.

Participants will be excluded for any of the following reasons:

- *Insufficient English Language Exposure.* Infants must be learning English as one of their primary languages. For inclusion in the final sample, at least 75% of children's language exposure must be English, as measured through parental report. Thus, multilingual infants will be included so long as they meet the minimum exposure criterion for English.
- *Pre-term.* Infants born 28 days or more before their due date will be excluded from the final sample.
- *Developmental Concerns.* Infants for whom caregivers report a developmental concern (e.g., uncorrected vision problems, Autism, or hearing concerns) or four or more ear infections in the last year will be excluded.
- *Data contribution.* An infant must contribute data for at least 50% of the trials (24 of 48 possible trials) after applying trial-level exclusion criteria (see below) across both sessions to be included in the final sample. Infants who only complete one session will still be included in the sample if they provide sufficient data from a single session (N = 24 trials).
- *Side bias.* We will analyze whether infants preferentially fixate a single side of the screen prior to the onset of the target word (baseline window: -1500 ms - 0 ms prior to target word onset). Infants who fixate one side of the screen for more than 80% of the baseline window on 80% or more trials in a given testing session will be excluded from further analysis (including data from their other test session).

- *Technical issues.* Infants will be excluded if the experiment logs or a comment from the caregiver indicates that there was a technical issue during the entirety of the study (e.g., if audio never played during the experiment) or if the video quality is too poor to allow for coding of infant looking (e.g., due to poor lighting conditions or poor video quality throughout).

S1.2. Trial-Level Exclusions.

A trial will be excluded for each of the following reasons:

- *insufficient looking data:* if the child is not fixating either object for more than 50% of the critical window (300ms-2800ms) or the baseline window (-2000ms - 0ms).
- *parental interference:* if the caregiver interferes with the task, e.g., if during the trial caregiver impedes the child's ability to fixate the screen, is talking to the infant during the critical window, or is pointing to images on the screen during the trial.
- *auditory interference:* if there is some form of auditory interference during the onset of the target word, including if the child is speaking during the onset of the target word.
- *technical error:* if the experiment logs or comments from the caregiver indicate that there was a technical issue during the trial that significantly impacts the design (e.g., audio did not play on a given trial or connectivity issues).
- *codability of the video:* if a given trial is impossible or extremely difficult to code (e.g., due to poor lighting conditions or poor video quality).
- *Low frame-rate trials.* Relatedly, trials will be excluded if the frame rate of a given trial is below 15 Hz on average (rounded up). Lower frame rates likely reflect low internet connectivity and introduce a higher level of imprecision in the planned analyses.

S2. Stimulus Characteristics

S2.1. Word Properties.

Table S1. Age of acquisition and word frequency for basic-level category terms

Word	Wordbank AOA (Comprehension, in months)	CHILDES log word frequency
bird	13	8.03
cat	12	8.43
dog	10	8.51
fish	15	7.96

Note: AOA is for American English. AOAs are the point at which 50% of infants are predicted to comprehend a given word, according to parental report. The prediction is based on fitting a logistic regression curve to each item's data in Wordbank, using the `fit_aoa()` function from the `wordbankr` package. The CHILDES log word frequency is based on the number of occurrences of each word in the American English corpus of CHILDES (Braginsky, Yurovsky, Marchman, & Frank, 2019; MacWhinney, 2000).

S2.2. Image Properties.

Table S2. Average typicality ratings for the image stimuli in each category.

Category	Typical items	Atypical Items
bird	M = 4.16, SD = 1.02	M = 2.78, SD = 1.26
cat	M = 4.29, SD = 0.88	M = 1.87, SD = 1.00
dog	M = 4.46, SD = 0.85	M = 2.24, SD = 1.25
fish	M = 4.12, SD = 1.04	M = 2.51, SD = 1.27

S3. Adult Norming Studies

We conducted two norming studies to assess the typicality and nameability of our image set (70 category exemplars across the four animal categories) during the stimulus selection process. Experiment scripts, stimuli, and norming data are publicly available on OSF (https://osf.io/3t8gf/?view_only=d4031533ac8549d5a8682b7819ba7bee).

Typicality Norming study. We collected data from 82 participants on Amazon Mechanical Turk. Participants were paid \$0.65 for participating in the study. Each participant judged the image sets from two out of 4 categories (*bird, cat, dog, fish*) in terms of their typicality. The image sets within a given category ranged from 16-20 unique images that varied in typicality. For each image, participants were asked “How typical is this X of [basic category label] in general?” and subsequently rated typicality on a Likert scale from 1 (very atypical) - 5 (very typical).

Nameability Norming study. To assess the nameability of each image (i.e., how consistently participants labeled a given image in our stimulus set in general, and how consistently the basic level category term was used), we conducted a norming study in which we collected data from 288 participants on Amazon Mechanical Turk. Participants were paid \$0.25 for participating in the study. Each participant judged 8 images, 2 from each category, presented in random order. For each image, participants were asked to name the image in 1-2 words. We included two catch trials presented randomly within the set of target images containing clearly identifiable objects (an image of a horse and a giraffe, respectively). Participants were only included in the final sample if they correctly named these items. 13 participants were excluded for failing to provide correct responses for the catch trials.

S4. Parent Survey Assessing Infant Experience-Based Typicality

Parents will complete a 24-item survey that corresponds to the 24 images that their infant saw during the looking-while-listening task. We will ask parents to rate the typicality of each item for their child on a five-point scale from “very atypical” to “very typical”. We will also query parents to provide additional information regarding the contexts in which their child experiences each item. We will provide the following instructions for the survey:

First, we will show you various images of animals and ask you to rate how typical the image is of the animals of this kind that your child experiences. To answer this question, please think about all of the media (books, TV, iPad, etc.) and real-world experiences (zoo, aquarium, etc.) your child has with animals. For example, you will be shown a series of birds and asked how typical each one is of the kinds of birds your child experiences. We will ask you to respond on a 5-point scale from "very atypical" to "very typical". As an example, an image of a sailboat that appears in a book that you read to your baby many times a week would be very typical of the kind of boats your child experiences.

S5. Video Processing and Coding

S5.1. Video Processing

Video data collected online via Lookit varies by participant and by trial depending on the local recording setup in the participant’s home (e.g., participants’ webcam, internet connectivity). In particular, the frame rate of participant videos can vary both across recording sessions and across trials within a specific recording session. This raises the issue of processing the data into a unified frame rate format. Our video coding and analysis pipeline approaches this problem as follows:

- Using Lookit software, we obtain a video recording for each trial in a given session for each participant. Lookit also stores the frame rate of the video for each individual trial along with its session data.
- For hand-coding individual videos, we concatenate all trial-level videos for a given session (viewed by a specific participant) into one session-level video, while retaining all individual frames (i.e. no data is lost in this process). Coders then code look location on a frame-by-frame basis using custom software (see the Video Coding section for further details). For the purposes of coding, we unify the frame rate across all trials in each session when concatenating the videos. This is necessary because the coding software used requires videos to be encoded in a consistent frame rate.
- After coding is complete, the trial time is subsequently converted back to the correct trial-level timing using the frame rate stored for a given session, ensuring that each frame is assigned the correct timing in the context of the trial.

S5.2. Video Coding

Children's looking behavior will be video recorded for each trial (using participants' local webcam in the case of online data collection). The eye movements will be coded frame-by-frame using custom software (Olson et al., 2020) by trained coders who are unaware of the target label, the trial-specific images, and the target location on each trial. On each frame, coders judge whether participants are looking left, right, or "off", i.e. not looking at either of the two images on screen. To determine reliability, 25% of the final sample will be independently re-coded. The reliability of three measures will be assessed: (a) the proportion of all frames on which coders agree on gaze location (frame agreement), (b) the proportion of gaze shifts on which coders

agree within one frame (shift agreement), and (c) the proportion of trials in which coders agree on the number of events coded (comparable trials).

Our current lab standards for reliability for in-person experiments are 97% for frame agreement and shift agreement, and 80% for comparable trials aggregated across all files that are recoded for reliability. In our initial online experiments, we have found it difficult to achieve the same reliability standards, due to the higher variability in recording quality and inconsistency in the positioning of the child relative to the screen and the experimenter view of the child's eye movements (for example, some children are positioned very close to the screen, while others are sitting farther away, impacting the ability to judge gaze shifts). In particular, variability in screen size and distance from the screen influences the child's visual angle and can make judging shifts between images and whether looks are on- or off-screen more difficult across set-ups. Based on our pilot (and other online experiments conducted in the lab), we are generally obtaining reliability measures of around 7 to 10% lower than the in-lab standards. We therefore aim to achieve minimum reliability levels of 90% agreement for frame agreement and shift agreement and 75% agreement for comparable trials in the main experiment. For any videos that fail to meet these reliability criteria, an independent third rater will resolve outstanding disagreements in coding.

S6. Effect sizes from previous studies

Table S3. Overview of effect sizes and analytic approach in past studies most closely related to the current work.

Study	Sample	Age (mo)	N	Question	Analytic Approach	Effect Size
Garrison et al. (2020)	younger infants	11.56 - 14.69	15	word recognition for familiar vs. unfamiliar exemplars	subject-level proportion target looking within window of analysis (367-2000 ms)	$d = 0.28$
	older infants	14.69 - 18.23	15			$d = 0.03$
Perry & Saffran (2017)	full sample	19.9 - 22.7	45	word recognition for exemplars with prototypical vs. atypical colors	subject-level proportion target looking within window of analysis used in the study (300-1800 ms), recomputed using publicly available data via Peekbank (Zettersten et al. 2021)	$d = 0.20$
Meints et al. (1999)	12-mo-olds	11 - 15	21	word recognition for typical vs. atypical exemplars	Increase in looking to the target from a pre-naming (2100 ms window prior to target onset) to a post-naming phase (3400 ms window after the target word onset). A direct comparison between typical and atypical items is not reported. The effect size is estimated as the effect size difference in above-chance accuracy for typical vs. atypical items. <u>12-mo-olds</u> : typical: $d = 0.77$; atypical: $d = 0.13$ <u>18-mo-olds</u> : typical: $d = 0.47$; atypical: $d = 0.99$ <u>24-mo-olds</u> : typical: $d = 0.56$; atypical: $d = 0.61$	$d \sim 0.64$
	18-mo-olds	16 - 21	38			$d \sim -0.52$
	24-mo-olds	22 - 25	23			$d \sim -0.06$
Southgate & Meints (2000)	full sample	18 - 24	54	preference for a typical vs. an atypical exemplar during word recognition	Increase in looking to the typical item (vs. the atypical item) from a pre-naming (2750ms window prior to target onset) to a post-naming phase (2800 ms window after the target word onset).	$d = 0.60$

Note. Effect sizes were computed based on test statistics reported in each paper or recomputed based on openly available data with the help of the effectsize package in R (Ben-Shachar et al. 2020).

S7. Online Pilot Experiment

S7.1. Online Pilot Experiment Methods

We conducted a pilot experiment on the Lookit platform to assess the feasibility of testing infant word recognition using the looking-while-listening procedure online and to inform decisions about age range in the current study. We tested infants between 12-18 months of age, collecting data from 51 unique participants in total. Families were paid \$5 for participating in the study. See S7.2 for a preliminary analysis of 41 participant sessions.

The experiment followed the same design and procedure as that outlined in the planned methods of the Stage 1 manuscript, with the following discrepancies:

- *Image stimuli.* Infants were tested on 16 different image stimuli, 2 typical and 2 atypical for each of the four categories (*bird, cat, dog, fish*). The images were a subset of the 24 images depicted in Figure 1 of the main manuscript.
- *Trial and session number.* Infants were tested in a single session with 16 test trials. Infants were randomly assigned to one of 8 trial orders (the specific orders can be viewed at the project OSF page).
- *Yoked categories.* We yoked the images from pairs of categories together in the pilot experiment. Specifically, dog images were always paired with bird images, and cat images were always paired with fish images.

S7.2. Online Pilot Experiment Results

51 12-18-month-olds were tested in the pilot experiment. Of these participants, 86% (44 of 51) contributed video data that could be coded after pre-screening. The six participant videos removed during pre-screening were excluded due to technical issues (5), parent interference (1), or infant fussiness/ insufficient looking data (1). Of the remaining 44 participants, we have thus

far coded 41 sessions. Among these 41 participants, 82.3% of trials contained sufficient looking data to meet our trial-level inclusion criteria (at least 50% looking during both the baseline window and the critical window; a comparable trial-level exclusion rate to in-lab studies with this age range, e.g., Garrison, 2020), and 37 of the 41 participants contributed sufficient looking data on at least half of the 16 experiment trials ($M = 12.7$ trials). Initial analyses using the analytic approach outlined in Aim 1 of the analysis plan (see below) suggest that infants were able to successfully recognize words across the experimental session. Overall, participants' average proportion looking to the target image during the critical window showed evidence of successful word recognition $M = 0.069$ (95% CI = [0.031, 0.106]; chance = 0; see Figure S1A). However, participants who were 14 months of age or older ($n = 21$) were notably more accurate ($M = 0.104$, 95% CI = [0.049, 0.159]; see Figure S1B and Figure S2) than participants below 14 months of age ($M = 0.023$, 95% CI = [-0.021, 0.067], $n=16$), consistent with observations in prior work (Bergelson, 2020). Word recognition was comparable for each target word and robust among infants older than 14 months of age, with somewhat lower accuracy for the target word *bird* (see Table S4). These initial results support the feasibility of using the Lookit platform to collect high-quality infant data using the current design and inform our decision to set the lowest age for our study at 14 months. They also suggest the presence of substantial variability, supporting our plans for individual differences analyses (Aim 3).

Figure S1. Timecourse of average proportion target looking for typical and atypical items in the Lookit-based pilot data (A) across all participants ($n=37$) and (B) for participants 14 months of age and older ($n=21$). Error bars are ± 1 SEs. The critical window extends from 300 ms to 2800 ms and the baseline window extends from -2000 ms to 0 ms relative to target word onset (highlighted in grey). Smoothed fits are based on a general additive model using cubic splines to visualize average looking trajectory.

Figure S2. Relationship between age (in months) and each participants' average baseline-corrected proportion target looking during the critical window. Each dot represents a participant. The blue line represents a linear regression fit. Error bars and error band represent 95% CIs.

Table S4. Overview of average baseline-corrected proportion target looking for each target word

Target Word	Infants <14 months (n = 16)	Infants > 14 months (n = 21)	Overall (n = 37)
bird	M = 0.0149 95% CI = [-0.053, 0.083]	M = 0.050 95% CI = [-0.073, 0.174]	M = 0.035 95% CI = [-0.038, 0.108]
cat	M = 0.029 95% CI = [-0.043, 0.102]	M = 0.111 95% CI = [0.009, 0.213]	M = 0.076 95% CI = [0.011, 0.140]
dog	M = 0.002 95% CI = [-0.114, 0.119]	M = 0.127 95% CI = [0.015, 0.239]	M = 0.073 95% CI = [-0.007, 0.153]
fish	M = 0.036 95% CI = [-0.044, 0.116]	M = 0.110 95% CI = [0.002, 0.217]	M = 0.078 95% CI = [0.010, 0.146]

References

- Ben-Shachar, M., Lüdtke, D., & Makowski, D. (2020). effectsize: Estimation of Effect Size Indices and Standardized Parameters. *Journal of Open Source Software*, 5(56), 2815.
<https://doi.org/10.21105/joss.02815>
- Bergelson, E. (2020). The comprehension boost in early word learning: Older infants are better learners. *Child Development Perspectives*, 14(3), 142–149.
<https://doi.org/10.1111/cdep.12373>
- Braginsky, M., Yurovsky, D., Marchman, V. A., & Frank, M. C. (2019). Consistency and variability in children’s word learning across languages. *Open Mind*, 3, 52–67.
https://doi.org/10.1162/opmi_a_00026
- Garrison, H., Baudet, G., Breitfeld, E., Aberman, A., & Bergelson, E. (2020). Familiarity plays a small role in noun comprehension at 12–18 months. *Infancy*, 25(4), 458–477.
<https://doi.org/10.1111/infa.12333>
- MacWhinney, B. (2000). *The CHILDES Project: Tools for Analyzing Talk*. Mahwah, NJ: Erlbaum.

- Meints, K., Plunkett, K., & Harris, P. L. (1999). When does an ostrich become a bird? The role of typicality in early word comprehension. *Developmental Psychology*, 35(4), 1072–1078. <https://doi.org/10.1037/0012-1649.35.4.1072>
- Olson, R. H., Pomper, R., Potter, C. E., Hay, J. F., Saffran, J. R., Ellis Weismer, S., & Lew-Williams, C. (2020). *Peyecoder: An open-source program for coding eye movements*. Zenodo. <https://doi.org/10.5281/zenodo.4313832>
- Perry, L. K., & Saffran, J. R. (2017). Is a pink cow still a cow? Individual differences in toddlers' vocabulary knowledge and lexical representations. *Cognitive Science*, 41(4), 1090–1105. <https://doi.org/10.1111/cogs.12370>
- Southgate, V., & Meints, K. (2000). Typicality, naming, and category membership in young children. *Cognitive Linguistics*, 11(1/2), 5–16. <https://doi.org/10.1515/cogl.2001.011>
- Zettersten, M., Bergey, C., Bhatt, N., Boyce, V., Braginsky, M., Carstensen, A., DeMayo, B., Kachergis, G., Lewis, M., Long, B., MacDonald, K., Mankewitz, J., Meylan, S. C., Saleh, A. N., Schneider, R. M., Tsui, A., Uner, S., Xu, T. L., Yurovsky, D., & Frank, M. C. (2021). Peekbank: Exploring children's word recognition through an open, large-scale repository for developmental eye-tracking data. *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*.