Cognition 193 (2019) 104033

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Crying helps, but being sad doesn't: Infants constrain nominal reference online using known verbs, but not known adjectives

Kristen Syrett[a,*], Alexander LaTourrette[b], Brock Ferguson[b], Sandra R. Waxman[b]

[a] Rutgers, The State University of New Jersey, New Brunswick, United States
[b] Northwestern University, United States

ARTICLE INFO

Keywords:
Word learning
Infants
Nouns
Verbs
Adjectives
Animacy

ABSTRACT

Speakers can make inferences about the meaning of new words appearing in an utterance based on the lexical semantics of other words that co-occur with them. Previous work has revealed that infants at 19 and 24 months of age can recruit the semantic selectional restrictions of known verbs (e.g., *eating*) to deduce that a noun appearing in the subject position maps onto an animate referent. We asked whether this ability to capitalize on the semantics of familiar words to identify the referent of a novel noun in subject position extends to adjectives, which also denote properties, and which also have animacy constraints (e.g., *hungry*). We found that unlike in the previous studies with verbs, neither 24- nor 36-month-olds could successfully recruit known adjectival semantics in an online task to home in on an animate nominal referent. However, 36-month-olds were successful in a more interactive, forced-choice version of the task without such strict time limitations. We discuss multiple non-mutually-exclusive hypotheses for this pattern of results, focusing on the role of the morphosyntactic cues, the (lack of) perceptual cues for the target property in context of the utterance, truth conditions, and cross-linguistic implications. These possibilities raise fundamental questions about the infant's developing lexicon and the linguistic and conceptual mechanisms at play in the process of word learning.

1. Introduction

If in the course of a conversation at a crowded event, you overheard the fragments of another speaker's utterance as, *"_ is eating the _"*, you could reasonably make certain inferences about the missing subject and object arguments of the verb *to eat*. For example, you could deduce that the object is something edible, and the subject is something or someone who can consume food, and is most likely an animal or human (absent a metaphorical usage). Thus, without even being presented with a visual scene or hearing the complete utterance, you are able to constrain potential meaning for the words that fill those argument slots and their corresponding referents. This process works because verbs such as the one featured here have semantic selectional restrictions on the arguments they take, and the semantics of verbs and the events they denote have consequences for the syntactic structure and other lexical items that appear in them (Chomsky, 1965; Grimshaw, 1979; Jackendoff, 1990, 1997; Pinker, 1994; Resnik, 1996; Van Valin, 1990). These inferences about permissible co-occurring meaning based on lexical semantics take many shapes. For example, only a liquid gets *sprayed* or *washes* something away, *wipe* requires one of the arguments to be a surface location, only a hard object *shatters*, only animate, sentient agents can *think* or *believe*, and so on (Levin & Rappaport Hovav, 1991; Rappaport Hovav & Levin, 2001).

However, the role of semantic content in constraining meaning goes beyond the selectional restrictions of verbs. For example, a word that co-occurs in the same utterance might also perform a similar function (Pinker, 1994). For example, a speaker who says, *Yum!* or *Mmm, delicious!* when offered a *larp* invites an inference from a listener that a larp is something edible (and comestibly desirable). Even a more general expression of surprise or one of clear positive or negative valence towards an object can help toddlers constrain the possible word meaning of a novel word and its potential referent (for review, see Tomasello, 2001). Thus, the lexical semantic representation of words can help to narrow the meaning of other words occurring within the same utterance (with the caveat that other operators such as negation or ironic prosody will lead to the opposite inferences).

1.1. Semantic restrictions in development

This ability of speakers to use known words to constrain the hypothesis space for upcoming linguistic information and their corresponding object referents has been well documented, even in children

* Corresponding author.
E-mail address: kristen.syrett@rutgers.edu (K. Syrett).

https://doi.org/10.1016/j.cognition.2019.104033
Received 5 November 2018; Received in revised form 27 July 2019; Accepted 28 July 2019
0010-0277/ © 2019 Elsevier B.V. All rights reserved.

as young as two years of age. In a seminal work that launched a series of eye gaze studies, Altmann and Kamide (1999) tested Altmann's (1999) hypothesis that a verb's semantic restrictions on its argument(s) are activated predictively and constrain the search for a discourse referent that satisfies those selectional restrictions. They showed adult participants a 'visual world' scene (Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995) with a boy, a ball, a cake, a toy car, and a toy train. Participants heard either *The boy will move the cake* or *The boy will eat the cake*. Anticipatory looks to the cake began after the onset of the verb *eat* (but not with *move*), indicating that participants rapidly integrated the lexical semantics of the verb and real world knowledge with the discourse context. Kamide, Altmann, and Haywood (2003) extended these findings to show that participants also rapidly integrate information about the agentive semantic role of the subject with the verb. (See Boland (2005) for findings related to the argument/adjunct distinction in prediction.)

A similar pattern of anticipatory looks reflecting integration of semantic constraints with the discourse context has been shown in much younger participants. Borovsky, Elman, and Fernald (2012) showed that even children as young as 3–10 years of age can make inferences about a possible object given a forced choice between four candidates in a visual world, based on the combination of the agent in subject position and the verb. Participants were shown a pirate ship, a cat, a bone, and a box of treasure. Looks to the treasure increased upon hearing *The pirate hides…*, whereas looks to the ship increased upon hearing *The pirate chases…* Patterns changed when the subject was *The dog*, and looks to the bone and the cat increased instead (see also, Borovsky and Creel (2014).).

Even younger children still make use of lexical information predictively in a scaled-down task similar to that of Altmann and Kamide's. Mani and Huettig (2012) presented two year olds with a picture of a cake and a picture of a bird, and presented them with an auditory stimulus with either an informative verb (e.g., *The boy eats the big cake*) or a neutral verb (e.g., *The boy sees the big cake*). Toddlers fixated on the target (the cake) soon after the onset of the verb *eat*, but waited until the onset of the noun with *see*, indicating that they made use of the verb's semantic restrictions to select a thematically appropriate argument referent. While the anticipatory abilities of toddlers—especially those with elevated vocabulary—had previously been documented in their rapid selection of an object referent for a familiar noun (Fernald, Zangl, Portillo, & Marchman, 2008), this finding demonstrated that they could predict a referent based on known selectional restrictions. (See Valian, Prasada, and Scarpa (2006) on consequences of verbal selectional restrictions on ease of production in two-year-olds.)

This finding can be extended to cases in which the object (or the nominal label) is not a familiar one. Goodman, McDonough, and Brown (2008) have shown that at two to 2.5 years of age, children can identify the correct animate referent when told to *Show me the ferret* after hearing a linguistic stimulus in which the selectional restrictions of the verb call for an argument that can consume food (*Mommy feeds the ferret*). However, this study does not tell us about the time course of their selection, and uses actual—albeit infrequent and potentially unfamiliar—nouns.

Remarkably, at the same age, two-year-olds can also use information about known nouns to make inferences about the selectional restrictions of novel verbs that take them as object arguments–even without a visual scene to support the inference. Yuan, Fisher, Kandhadai, and Fernald (2011) presented children with a series of linguistic utterances during a listening phase in which one novel verb was paired with inanimate objects (e.g., *She nerked a fork*), while another was paired with animate arguments (e.g., *She wants to stipe the elephant*). During the test phase, participants were presented with a choice between an animate and an inanimate candidate (e.g., a cat and a spoon), and asked which one they would *nerk/stipe*. In both these unfamiliar trials and in trials with familiar verbs (e.g., *eat, drive*), participants zeroed in on the intended referent after the onset of the verb.

1.2. Narrowing the space for subject arguments

If the selectional restrictions of a verb can be rapidly recruited to identify an intended referent for an object argument, then they may also be useful in narrowing down the potential meaning of a subject argument. Indeed, in a clear illustration of young children's capacities to recruit known word meanings to deduce the meaning of novel words, Ferguson, Graf, and Waxman (2014) have shown that by 19 months of age, infants can use the animacy restrictions of known verbs to constrain the meaning of novel nouns in subject position. Ferguson et al. (2014) compared infants' eye gaze across two trial types: one in which the objects seen were familiar and referred to with known nouns (e.g., *dog, bottle*), and one in which the objects were unfamiliar and referred to with novel nouns (e.g., *larp*). Within these critical unfamiliar trials, there were two further conditions: a neutral condition in which the property made no reference to animacy (e.g., *The larp is right here.*) and an *informative* condition in which the verb selects for an animate subject (e.g., *The larp is crying*). After a brief exposure to side-by-side images of an animate and an inanimate object, the key linguistic information was presented during a trial phase when there was no supporting visual scene. In the next phase, infants saw the same two images again and were asked where the larp was, and their fixations to one or the other image during this Test phase was measured. Ferguson et al. (2014) found that 19-month-old infants quickly and successfully identified the correct referents in the familiar trials, and while infants had a slight overall preference for animate referents in the unfamiliar trials, their looks were mediated by the type of predicate: novel nouns that were combined with animate-subject-selecting verbs yielded more looks to the animal than the artefact at test.

Ferguson, Graf, and Waxman (2018) extended this research to 24-month-olds, replicating the same pattern observed with the 19-month-olds. Both groups of infants successfully recruited the known verbal meanings to zero in on an animate referent in the unfamiliar trials. The main difference Ferguson et al. (2018) found was that the 24-month-olds were more efficient in their processing, and therefore settled in on looking at the intended referent much sooner (approximately 2200 ms sooner). Thus by two years of age, infants not only recruit this semantic knowledge about animacy in the service of referent selection, but do so quickly, as the speech stream unfolds.

These combined findings illustrate the power of the developing lexicon to rapidly constrain the meaning of newly encountered words, and in particular, highlight the potency of the selectional restrictions of verbs in narrowing the hypothesis space for the meaning of nouns appearing in their argument positions. Given the potency of verbs in constraining the space of an intended object referent for a novel noun, a key question that arises is how general this process of recruiting known words to learn new ones is, and whether a lexical item from another category with similar animacy restrictions can perform the same function.

1.3. Beyond verbs

To begin to investigate whether these previous results generalize across grammatical categories, we turn to adjectives as a potential cue. Here, we outline four main reasons for doing so. First, verbs and adjectives share a similar semantic representation in that both verbs and adjectives denote *properties*. While verbs may typically be seen as denoting properties of events and adjectives properties of entities, this distinction is misleading. While adjectives do consistently modify nouns in the syntax, these nouns may in some cases refer to events (e.g., *performance, war, nap, party*) (Grimshaw, 1990; Zucchi, 2013), even with nouns that seemingly only refer to individuals (e.g., *I began the book/the book was long*, Pustejovsky, 1995). Nor are adjectives restricted to describing properties of *individuals*, since they can also describe *sets* of individuals (e.g., *the group was large, the problems were numerous*, etc.). Here, too, there is a parallel between adjectives and verbs, since

some verbs can or must apply to a group or plurality (e.g., *the crowd dispersed, the committee disbanded, the children gathered around the teacher*). With these particular lexical items, there is a shared conceptual and semantic interpretation regarding the truth conditions that must hold in order for the predicate (verb or adjective) to be true, and in each of these cases, outside of a metaphorical usage, the properties share a restriction on the agent expressed by the grammatical subject based on animacy (e.g., an agent who is *crying* or is *sad* must have a particular biological constitution and/or be able to experience emotion, an agent who is *eating* or is *hungry* is consuming or desires to consume food, etc.).

Second, verbs and adjectives have similar surface-level distribution (at least in English, the language under investigation in Ferguson et al. (2014, 2018) and here). Both can appear in an utterance-final position after *is*. The verb can appear in a present progressive form (i.e., *X is eating*) while the adjective can appear in post-copular predicative position (i.e., *X is hungry*). This similarity allows us to construct the relevant linguistic stimuli in a way that is minimally different. Third, both verbs and adjectives appear around the same time in child-directed speech (Dale & Fenson, 1996), and are typically delayed in production relative to nouns. Thus, we can be confident that neither category is significantly privileged relative to the other as, for example, nouns would be to either.

Finally, much like verbs, adjectives can also help children to rapidly home in on an intended object referent in a visual scene (Eberhard, Spivey-Knowlton, Sedivy, & Tanenhaus, 1995; Sedivy, Tanenhaus, Chambers, & Carlson, 1999; Tanenhaus et al., 1995). This ability to recruit adjectives incrementally in the service of narrowing down the hypothesis space surfaces in children as early as three years of age (Fernald, Thorpe, & Marchman, 2010; Thorpe, Baumgartner, & Fernald, 2006). For example, when presented with a forced choice between a blue car and a blue house and asked, *Where's the blue car?*, 36-month-olds waited until soon after the onset of the noun to lock in on looking at a referent, but when shown a blue and a red car, and asked the same question, they did not wait, and instead made immediate use of the adjectival modifier. Children only six months younger were (rather counterintuitively) less efficient in this latter condition, and did not show signs of rapid integration of the adjectival meaning.

Under one hypothesis (Ninio, 2004; Thorpe et al., 2006), integration of the prenominal adjective with the following noun with which it semantically composes and syntactically combines is challenging for children younger than age three. Their inferior processing capacity might mean that they cannot encode the linguistic information in light of the visual discourse scene, and also attend to the subsequent linguistic information that follows in the speech signal, which impacts the representation they are able to build incrementally. If this is the case, then we might predict that removing these barriers paves the way to successful recruitment of adjectival information in children younger than 36 months of age. Specifically, we could take two steps.

First, we can remove the adjective from prenominal position and place it in salient predicative position, which also does not carry restrictions on which adjectives can occur in that position, as the prenominal position does (Prasada, 1992). In this way, too, the stimuli would parallel previous stimuli with verbs. Second, we can follow the previous research with verbs and disentangle the linguistic and visual stimuli by presenting the relevant linguistic information in the absence of a visual scene—a strategy that has been successfully used in a number of word learning studies (Arunachalam & Waxman, 2010; Arunachalam, Syrett, & Chen, 2016; Yuan & Fisher, 2009). One might also hypothesize that the 30-month-olds struggled in Fernald et al's task, because they were confronted with color words, a subcategory of adjectives for which development is notoriously prolonged. (For discussion of the challenges of learning color words, see in particular Sandhofer and Smith (1999), Soja (1994), and Wagner, Jergens, and Barner (2018).) Thus, by appealing to adjectives that have animacy restrictions (as the verbs that were previously tested have), we not only conduct a study that parallels previous work with animacy-selecting verbs, but we also sidestep difficulties associated with other types of adjectives.

The hypothesis that children might recruit familiar adjectives to assign a meaning to novel nouns is bolstered by the fact that by 30 months of age, children are able to accomplish a task that seems equally hard if not harder: using a familiar adverb to deduce the meaning of other novel words. They recruit the presence of a highly frequent and familiar manner adverb (*slowly*) to assign a meaning to a novel word appearing earlier in the utterance (*He's gonna pilk it slowly*) as an event-denoting verb (Syrett, Arunachalam, & Waxman, 2014), instead of a referent-denoting noun, as reported in Arunachalam and Waxman (2010) with the visual scenes and same sentences in the absence of this adverb. They also recruit the presence of a known adverb (*together*) to narrow the hypothesis space of an unknown verb appearing with a coordinated subject (*My sister and the lady are going to biff together*) as referring to an event in which two agents are engaged in spatiotemporally contiguous subevents, rather than a causative event in which an agent acts on a patient (Arunachalam et al., 2016). Finally, children at this age are capable of using the selectional restrictions of known adverbs to constrain the meaning of novel adjectives in predicate position with which they compose (*These are both completely pelgy*), given a forced choice between objects displaying two competing properties corresponding to different gradable adjective scales (e.g., height– *tall* and transparency– *clear*) (Syrett & Lidz, 2010; Syrett, 2007).

1.4. Current research

We have reviewed evidence from previous work that by 19–24 months of age, children use the animacy restrictions of a known verb (even without a co-occurring supporting visual scene) to deduce the meaning of its subject argument and its corresponding real world referent. We also presented reasons for thinking that known adjectives with similar selectional restrictions and salient utterance position might be able to accomplish the same goal. Thus the question that we entertain here is whether adjectives possess the same power as verbs in word learning: can infants recruit the lexical-semantic representations of known adjectives to narrow the hypothesis space of a novel noun meaning, and zero in on the intended subject referent?

We began with the oldest age at which Ferguson et al. (2018) found that infants can use the animacy selectional restrictions of verbs to constrain the meaning of novel nouns: 24-month-olds (5 months older than the infants in their previous study). We conservatively started at this age instead of at the younger age of 19 months, to increase the chance of success. (As we will show, children do not succeed at 24 months of age.) Based on these findings, we then moved to an older age group, 36-month-olds, since children at this age not only produce a range of adjectives, but have been shown to display semantic knowledge of different adjectives (see Syrett, 2007), and can process known adjectives rapidly to identify a referent for the noun they modify (Fernald et al., 2010). Finally, given the findings with this older group, we tested this same age range on a more interactive forced-choice pointing task exploring the same semantic constraints.

Previewing our findings, we find that neither 24- nor 36-month-old children homed in on the animacy-restricted meaning of the novel noun/referent based on the presence of a known adjective in an online task, whereas they had done so efficiently in previous studies with familiar verbs. However, the older age group succeeds when the task is revised to an offline forced-choice pointing paradigm. These admittedly puzzling findings lead us to consider the possible reasons for the inability of adjectives to perform a role comparable to verbs in the online paradigm, and the possible reasons for the difference in success rate for the two versions of the task with adjectives, targeting linguistic, perceptual, and conceptual components as possible (not mutually exclusive) causes. As a consequence, these findings raise intriguing and deep questions about the informational sources and deductive processes at play in early word learning, and lay out a clear set of hypotheses for

further research in verbs and adjectives in English and cross-linguistically.

2. Experiment 1

2.1. Participants

45 24-month-olds were included in the final sample (M: 24.2 months, SD: 0.64, 23 females, 22 males). An additional 14 infants were excluded due to failure to complete the task (7), severe track-loss (failing to provide 25% looking on 3 or more familiar and unfamiliar trials: 5), or technical issues (failure to calibrate, 1, equipment malfunction, 1). All infants were recruited from Evanston, IL, based on caregiver reports of less than 25% exposure to a language other than English. Upon entering the laboratory, caregivers completed the MacArthur Short Form Vocabulary Checklist: Level II (Form A) (Fenson et al., 2000), as well as a supplementary checklist that asked which of the familiar adjectives used in this design were known by their child. Parental reports indicated that infants produced 59.4 words on the MCDI (SD = 23.2), including 34.8 (SD = 12.3) of 51 nouns and 5.8 (SD = 4.1) of 15 verbs. In addition, children knew an average of 4.50 of the 6 target adjectives used in the task,[1] and the majority of children tested were reported to know each of the individual adjectives, with the exception of *friendly*. The results of the MCDI scores strongly correlate with the results of the supplemental adjective checklist ($r(42) = 0.44$, $p = 0.003$). All experiments and the norming study reported below were performed under IRB approval, and informed consent was obtained for experimentation with human subjects.

2.2. Apparatus

We used a Tobii T60XL corneal-reflection eyetracker for stimulus presentation and data collection in a laboratory setting. This eyetracker has a sampling rate of 60 Hz, and a display size of 57.3 × 45 cm. Participants were seated approximately 60 cm from the screen on their caregiver's lap. Caregivers were given opaque glasses to wear so they could not see the stimuli.

2.3. Materials

2.3.1. Stimuli

There were 12 trials in each experimental session. Each trial had the same structure.

The trial consisted of three phases: a Preview phase, a Dialogue, and a Test phase, as depicted in Fig. 1.

2.3.1.1. Visual stimuli. Selection of familiar objects was evenly split between animate and inanimate objects and was based on names of objects understood by at least 72% of 15-month-olds (inanimate nouns: *bottle, spoon, car*; animate nouns: *bird, cow, dog*) (Dale & Fenson, 1996). Unfamiliar objects were those whose names infants would most likely not know (abstract sculptured artefacts and exotic animals). The visual stimuli were identical to those employed in Ferguson, Graf, and Waxman (2014, 2018), allowing us to make the comparison between familiar verbs and familiar adjectives as tight as possible and replicate previous baseline results. There was one minor exception in that we replaced an image of a horse with an image of a car during one of the familiar (non-target) trials in order to balance animacy among images. This did not affect the results for the familiar trials.

[1] In comparison, in Ferguson et al. (2018), 19-month-old infants knew on average 4.9 of the 6 'known' target verbs, and at 24 months, they were reported to knew 5.96 of the 6 target verbs. Both age groups succeeded in this task with verbs.

2.3.1.2. Linguistic and auditory stimuli. Auditory stimuli were modeled after the stimuli featured in Ferguson, Graf, and Waxman (2014, 2018), and were recorded in a sound-attenuated room by two native speakers of English (male, female) using a child-directed speech style.

Lexical items were selected based on a combination of factors in order to identify an optimal set of adjectives. First, we began with the animacy-selecting verbs from the previous studies by Ferguson, Graf, and Waxman (2014, 2018): *crying, dancing, drinking, eating, looking, sleeping*. We then identified corresponding adjectives for as many of these as possible: *sad* (*crying*), *hungry* (*drinking, eating*), *sleepy* (*sleeping*), and then supplemented the list with three more similar animacy-restricting adjectives: *angry, friendly, happy*. These adjectives are mono- or bisyllabic and are highly frequent. With the exception of *sad*, all had a salient and typical adjectival *–y* ending.

To confirm that these six adjectives had an animacy bias, we administered an independent norming task to 60 undergraduates (30 in each condition). We pseudo-randomized these six adjectives along with four other animate candidates (e.g., *mad, tired*), 11 inanimate candidates (e.g., *wooden, absorbent, plastic*), and 15 potentially neutral candidates (e.g., *nice, pretty, fuzzy, fake, wet*), and provided adults with a rating task, the instructions for which are in the Appendix. Participants were randomly assigned to either the *animate* or *inanimate* version of the task, with conditions evenly balanced. Participants were either asked to rate adjectives as animate (1: cannot be used to describe animate things…5: is ok being used to describe animate things) or inanimate (1: cannot be used to describe inanimate things…5: is ok being used to describe inanimate things).

The six target animate adjectives received an average of 4.97–5 for the animate version of the rating task and an average of 1.27–1.63 for the inanimate version, indicating that participants viewed them as selecting for 'animate things' and *not* 'inanimate things'. The target inanimate controls displayed the opposite pattern across the two rating tasks, averaging an overall 2.16 in the animate version and 4.63 in the inanimate version, indicating that participants viewed them as selecting for 'inanimate things' and not 'animate things'. We were therefore confident about our six target adjectives being robust cues for animacy.

The predicates in the *neutral informative* condition were the same as in Ferguson, Graf, and Waxman (2014, 2018): *right here, very clean, nearby, so close, so nice, so little*. None of these is a strong signal for animacy or inanimacy. The novel nouns were also the same as those in these previous studies. They were all monosyllabic and were CVC or CVCC words that are phonotactically licensed in English.

2.4. Procedure

There were 6 *familiar* trials and 6 *unfamiliar* trials. In the familiar trials, the animal was a common animal (e.g., dog), and the object was a known, familiar object (e.g., bottle). In the unfamiliar trials, the animal was an uncommon, unfamiliar animal and the object was a novel artefact. The position of the two images was counterbalanced across trials.

Each trial type had three phases. In the Preview phase (6 s), infants saw images of two objects (one animal and one novel object) presented side-by-side on the screen. In the Dialogue phase (9 s), an abstract screensaver appeared and infants heard a dialogue. In the Test phase (6 s), the same two images from the Preview phase then reappeared in the same positions on the screen.

During the Preview phase, infants heard either "*Look here!*" or "*Look at this!*" to direct their attention to the screen. The animal and object were displayed for 6 s, then disappeared. The trial then proceeded to the Dialogue phase. During this phase, as infants viewed a colorful, abstract scene, a dialogue between two people played. In the familiar trials, the speakers made reference to one of the two objects by making a statement about it with a full definite DP in subject position and a locative or adjectival phrase without animacy restrictions in predicative position (e.g., *The bottle is right here.*). In the unfamiliar trials, the

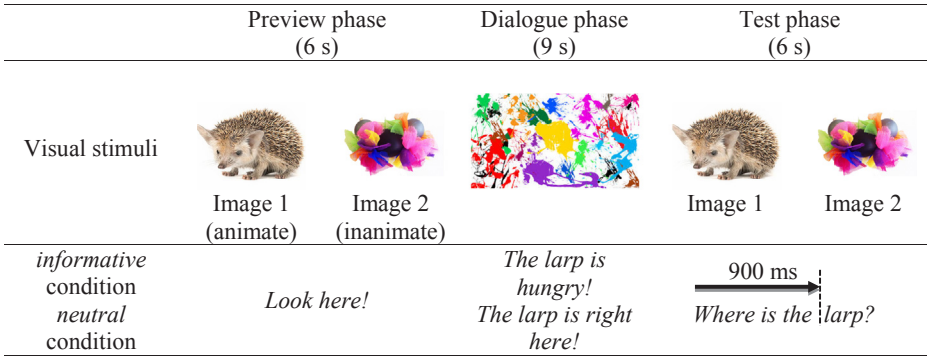

Fig. 1. Trial structure, with example of an unfamiliar trial in both the *informative* and *neutral* conditions.

speakers instead used novel object names (e.g., *The larp*). The unfamiliar trials were the locus of the experimental manipulation.

Infants were randomly assigned to one of two between-subject conditions (*informative* and *neutral*), depending on the adjective they heard within the unfamiliar trials during this phase. In the informative condition, the novel noun was accompanied by a familiar, highly frequent adjective, which seeks out an animate subject (e.g., *The larp is hungry*.). These utterances all featured similar intonation. In the neutral condition, the novel noun was accompanied by a locative or adjectival phrase that had no animacy restrictions, as in the familiar trials (e.g., *The larp is right here*.). In the next phase, the Test phase, the two candidate images from the Preview phase reappeared in their original locations, and infants were asked about the location of the object that had been mentioned during the preceding dialogue (e.g., *Where is the dog/bottle/larp?*). The onset of the target noun occurred approximately 900 ms into this phase.

To be clear, across both conditions, infants saw the same images and heard the same linguistic stimuli during the Preview and Test phases. The only difference between the conditions was in the linguistic information provided during unfamiliar trials of the Dialogue phase, as infants viewed the abstract screen saver. Thus, the question was whether this minimal linguistic information in the absence of any potential co-occurring object referents or event, could constrain infants' selection of a previously seen object referent in the Test phase.

2.5. Predictions

Based on the previous findings, we generated two main sets of predictions. First, we predicted that in the familiar trials, infants would rapidly and successfully appeal to these known nouns and the familiar images (e.g., *bottle*, *dog*) to zero in on looking at the correct referent during the Test phase. Second, we predicted that in the unfamiliar trials, where a novel noun was mentioned during the Dialogue phase, infants in the neutral condition would look at chance between the two candidates, since the linguistic stimulus was uninformative about the choice of referent. There was nothing in the Dialogue phase that could be used to direct infants' gaze towards one image or another at Test. By contrast, in the informative condition, infants who recognize the animacy-based selectional restrictions of the target adjectives (e.g., *sad*, *hungry*) should zero in on the animate nominal referent, despite it not being a familiar animal. Thus, we predicted that although infants might perhaps demonstrate an overall preference for the animals over the artefacts, this preference (both in terms of what was preferred at Test and how quickly infants locked in on this image) would be modulated by the linguistic information presented during the Dialogue phase.

2.6. Results

2.6.1. Data preparation and analysis

We analyzed participants' looking during a time window in the Test phase of each trial from 500 ms prior to noun onset until 4000 ms after noun onset. We chose this window to perform a comparable analysis with the previous research, which suggests the effects should emerge within this period, occurring as late as 3750 ms in the previous research with 19-month-olds (Ferguson et al., 2014, 2018). During this window, in the familiar trials, infants can rely on the presentation of images in the Preview phase and the familiar noun mentioned in the Dialogue phase to anticipate the referent being highlighted in the Test phase. Our dependent variable was the proportion of looking time devoted to the animate referent out of the total time spent looking at both the animal and artefact combined during this target window of the Test phase. The raw data for this and subsequent experiments are available in the Supplementary Materials.

To assess children's looking behavior, we created areas of interest (AOIs), measuring 810 × 710 pixels, around each of the test images. Any looking outside of the AOIs was excluded from the analysis. We also excluded any trial during which the child attended to the images for less than 25% of the relevant window at test. For inclusion in the analysis, children were required to contribute a minimum of 3 trials in both familiar and unfamiliar blocks. On average, children contributed an average of 5.80 trials (SD = 0.59) during the familiar trials and 5.51 trials (SD = 0.84) during the unfamiliar trials. Critically, the number of unfamiliar trials contributed did not vary by condition ($M_{informative}$ = 5.50, $M_{neutral}$ = 5.52, $t(43) = 0.09$, $p > 0.9$).

All data preparation and analyses were performed using R (R Development Core Team, 2012) with the eyetrackingR package (Dink & Ferguson, 2015), and followed the method employed in Ferguson et al. (2018). To analyze children's looking behavior during the target window, we conducted two kinds of analyses: an aggregate analysis and a time-course analysis. In the *aggregate* analysis, we averaged participants' looking across the full test window, finding the proportion of looking to the animate referent during each trial for each subject. To correct for the association between means and variances in proportional data, we submitted these proportions to an empirical logit transformation. We then created a linear mixed-effects model including random effects of participant and trial and random slopes when supported by the data and design (Bates, Kliegl, Vasishth, & Baayen, 2015); transformed proportions remained the dependent variable. Fixed effects (e.g., of condition) were dummy coded. Significance values were obtained using the Satterthwaite method for estimating degrees of freedom in the lmerTest package (Kuznetsova, Brockhoff, & Christensen, 2017), an appropriate method of significance evaluation for smaller samples. In all studies, traditional ANOVA methods yielded the same results.

For the *time-course* analyses, we separated the data into 100 ms bins and conducted a cluster-based permutation analysis across these bins (Maris & Oostenveld, 2007), again using logit-transformed proportions as the dependent variable. For this analysis, we first selected a t-threshold, based on a desired alpha of 0.05. Next, we conducted t-tests within each time-bin and identified the time-bins whose t-statistic

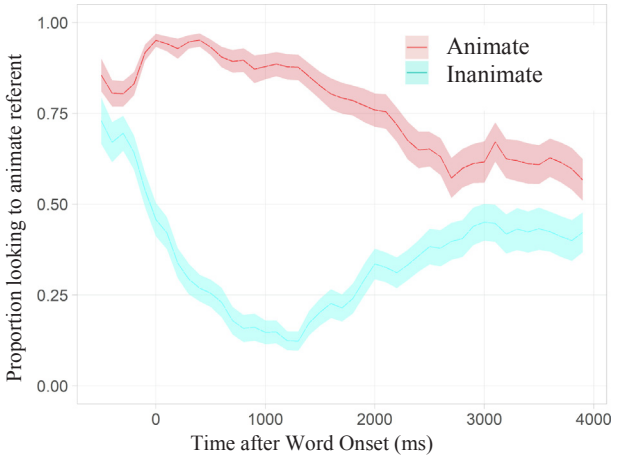

Fig. 2. Time-course of looking to the animate referent for the animate and inanimate target across the familiar trials of Experiment 1 (24-month-olds). Shaded regions represent ± 1 SEM.

exceeded our threshold. Finally, we summed together the t-statistics of any consecutive time-bins, creating a measure of the size of this divergence between groups. To evaluate the likelihood of such a divergence occurring by chance, we conducted 10,000 simulations of the data with randomly shuffled condition labels and performed the same analysis on each simulated dataset. To obtain a significance value, we then compare the observed divergence to this distribution of chance-based divergences. Summed t-statistics are reported as a measure of the effect size, but all significance values are derived directly from the permutation analysis.

2.6.2. Experimental results

For the two different trial types (Familiar, Unfamiliar), we begin by reporting the results for the aggregate analysis, then turn to the results for the time-course analysis. The results for the timecourse analysis are presented in Figs. 2 and 3. Fig. 2 presents the looks to the animate referent for the animate and inanimate targets in the familiar trials, while Fig. 3 presents the looks to the animate referent for the informative and neutral conditions of the unfamiliar trials.

2.6.2.1. Familiar trials. As predicted, children quickly and accurately identified the referent for the target noun in the Test phase of the familiar trials. Their looking during the these trials varied consistently as a function of the known words they heard: children looked significantly more to the animate object when hearing words referring to the animate object ($M = 0.79$, $SD = 0.12$) than when hearing words referring to the inanimate object ($M = 0.34$, $SD = 0.16$), $\beta = 3.24$, $SE = 0.34$, $p < 0.0001$. Moreover, this difference between animate-referring and inanimate-referring words emerged quickly: a cluster-based permutation analysis identifies a *significant and prolonged divergence* between the two trial types beginning 200 ms before the onset of the noun, extending 2900 ms after the onset (cumulative t-statistic = 303.9, $p < 0.0001$). This early onset of the first divergence is likely due to the structure of the task and the seconds leading up to the target prompt: during the Dialogue phase, the familiar referent was mentioned, and then participants were then immediately directed locate it in the Test phase.

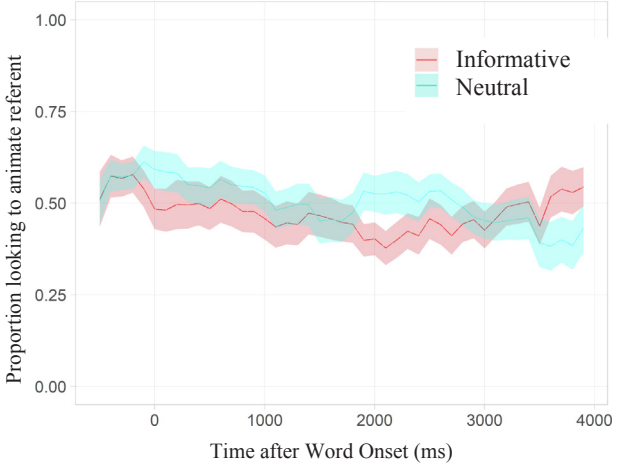

Fig. 3. Time-course of looking to the animate referent in the two conditions of the unfamiliar trials of Experiment 1 (24-month-olds). Shaded regions represent ± 1 SEM.

2.6.2.2. Unfamiliar trials. In contrast to the familiar trials, children in both the informative ($M = 0.47$, $SD = 0.12$) and neutral ($M = 0.51$, $SD = 0.10$) conditions did not significantly differ in their aggregate looking to the animate referent ($\beta = -0.22$, $SE = 0.27$, $p = 0.43$). A cluster-based permutation analysis found no significant divergences between the two conditions: no time-bin even exceeded the initial t-threshold. This similarity between these two conditions suggests that 24-month-old children were unable to use animacy restrictions on familiar adjectives to identify the referent of the novel nouns, in contrast to the previous performance with verbs.

Indeed, as a stringent test of this apparent null effect, we directly compared children's behavior here to the performance of the children in Ferguson et al. (2018) original study. Children were the same age across studies, $t(91) = 0.22$, $p = 0.83$, and viewed identical stimuli on the unfamiliar trials, making the results straightforwardly comparable. We began by constructing a linear mixed-effects model with the aggregate proportion of looking to the animate referent as the dependent variable and random effects of participant and trial (as above), but now, we included both Study (Verbs vs. Adjectives) and Condition (Informative vs. Neutral) as fixed effects. Results revealed significant main effects of Study, ($\beta = 0.20$, $SE = 0.043$, $p < 0.001$) and Condition ($\beta = 0.11$, $SE = 0.042$, $p = 0.013$), but these effects were qualified by a significant interaction between Study and Condition ($\beta = -0.15$, $SE = 0.06$, $p = 0.017$). As predicted, this interaction is driven by children's increased looking to the animate referent in Ferguson et al.'s Informative condition featuring verbs ($M = 0.65$, $SD = 0.17$) but not in our own Informative condition featuring adjectives ($M = 0.47$, $SD = 0.12$). To expand on this result, we also compared the time-course of looking between Ferguson et al.'s Informative (Verb) condition and our own Informative (Adjective) condition. A cluster-based permutation analysis revealed an extended divergence between the two conditions from 0 to 2600 ms after word onset. A similar analysis comparing Neutral conditions across studies yielded no significant divergences, $p > 0.5$. Thus, 2-year-olds were significantly better at using animacy-selective verbs than animacy-selective adjectives to infer novel referents.

2.7. Discussion

In Experiment 1, we investigated whether 24-month-olds could recruit the animacy constraints of familiar adjectives to identify the intended referent of a novel noun in the subject position. Previous research had documented that infants at the very same age with the very same paradigm were able to successfully use familiar verbs to perform this function. In contrast to the previous findings, the present findings indicate that 24-month-olds do not meet with success when presented with familiar adjectives. Nor did overall vocabulary knowledge or specific knowledge of the adjectives significantly predict success on the task. Given that this experiment gives us a minimal age at which infants cannot recruit adjectives for this purpose, we decided to conservatively target a much older group (36-month-olds) in the next experiment. This

age was selected based on the findings discussed earlier related to the processing capacity and semantic knowledge of adjectives possessed by children at this very age.

3. Experiment 2

3.1. Participants

Sixty-four 36-month-olds were included in the final sample (M: 35.8, SD: 1.3, 32 females, 32 males). Because data from 36-month-olds may no longer be directly comparable to the previous work with 24-month-olds, we conducted a power analysis of Ferguson et al. (2018) aggregate effect ($d = 0.67$), yielding our current sample size of 32 children per condition for 75% power. This sample size was also determined based on a power analysis of a previous sample size of 48 participants, which suggested that our original sample was underpowered. We therefore ran an additional 16 subjects, and observed the same set of results reported in 3.4. An additional 32 children were excluded due to failure to complete the task (7), severe trackloss (17), parental interference (1), or technical issues (failure to calibrate, 4, or equipment malfunction, (3). (In addition, the majority of children tested were reported to know each of the individual adjectives.)

3.2. Stimuli and procedure

The stimuli and procedure were the same as in Experiment 1.

3.3. Predictions

We predicted that the 36-month-olds would successfully home in on an animate referent, given the selectional restrictions of the known adjectives in this experiment.

3.4. Results

Data preparation and analysis were conducted as before. As in Experiment 1, we present the proportion of looking time in the two trial types first, followed by discussion of the time-course of looking within the two conditions of the two trial types (Figs. 4 and 5). Data preparation was identical to Experiment 1. Children contributed an average of 5.60 trials ($SD = 0.83$) in the familiar trials and 5.44 trials ($SD = 0.89$) in the unfamiliar trials. Again, the number of unfamiliar trials contributed did not differ in the informative ($M = 5.56$, $SD = 0.76$) and neutral ($M = 5.22$, $SD = 1.04$) conditions ($t(62) = 1.51$, $p = 0.13$).

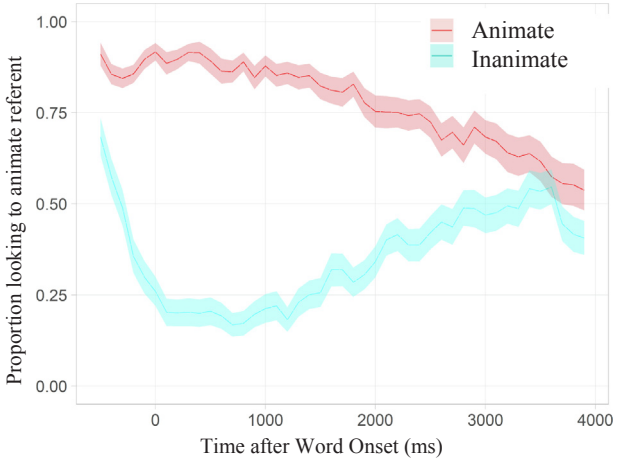

Fig. 4. Time-course of looking to the animate referent for the animate and inanimate target across the familiar trials of Experiment 2 (36-month-olds). Shaded regions represent ± 1 SEM.

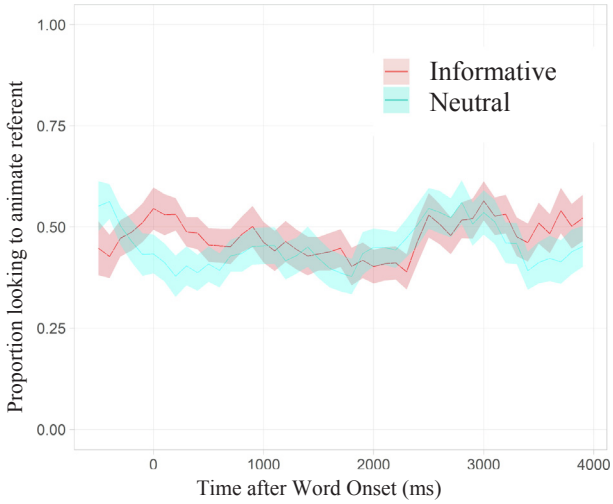

Fig. 5. Time-course of looking to the animate referent in the two conditions of the unfamiliar trials of Experiment 2 (36-month-olds). Shaded regions represent ± 1 SEM.

3.4.1. Familiar trials

Performance on the familiar trials reveals that the paradigm was also viable for 36-month-olds. On trials where the word referred to an animate object, children looked significantly more to the animate referent ($M = 0.80$, $SD = 0.12$) than on trials for which the word referred to the inanimate object ($M = 0.34$, $SD = 0.15$), $\beta = 3.31$, $SE = 0.26$, $p < 0.001$. Again, we observed a significant divergence between the two types of trials over time, $t_{summed} = 330.0$, $p < 0.0001$, lasting from 500 ms before word onset to 3200 ms after word onset.

3.4.2. Unfamiliar trials

In the informative condition, children did look slightly more to the animate referent ($M = 0.47$, $SD = 0.12$) than their counterparts in the neutral condition ($M = 0.43$, $SD = 0.11$); however, this difference did not reach significance ($\beta = 0.21$, $SE = 0.21$, $p = 0.31$). A cluster-based permutation analysis revealed no significant divergence between the two conditions, $ps > 0.4$, suggesting children in the informative condition failed to identify the animate target as the novel noun's referent.

3.5. Discussion

The results of this experiment indicated that 36-month-olds were able to rapidly home in on the object referent in the familiar trials, demonstrating that they performed successfully in the paradigm when not called upon to learn the meaning of a novel word. However, in the unfamiliar trials, children failed to home in on the animate referent. Thus, in contrast to the findings of the previous research, where infants were able at 19–24 months of age to recruit familiar verbs to identify an intended object referent, the current results show that children were unable to do so with familiar adjectives as late as 36 months of age.

Why should 36-month-olds, who have superior lexical knowledge, productive capacity, and processing efficiency, not exhibit similar success? We might entertain two possibilities. One is that the time pressure of the timed looking task impeded stable selection of the referent. This might happen, for example, if these older children were entertaining other options for the meaning of the novel noun, but were unable to resolve the indeterminacy before the end of the target utterance. The other is that the children had made a selection early in the test phase but then rapidly lost attention. For example, 36-month-olds might be so efficient in identifying the referent that they do so rapidly, and then, because the images themselves are not novel (having appeared earlier in the trial), and are static, they do not linger on their selection. In either case, we reasoned that correct selection of the

animate referent might be facilitated and detected in a more interactive task where the time limit was removed. We therefore adapted our paradigm as an offline forced-choice study, which we report in Experiment 3.

4. Experiment 3

4.1. Participants

22 children (M: 36; 2, SD: 1.27, range 34; 7–39; 3; 9 females, 13 males) participated.

4.2. Stimuli

We adapted the stimuli from the two previous experiments to create a forced-choice pointing task administered via Powerpoint slides on a 17″ Macbook Pro screen. The two images shown during the test phase were placed side by side on the screen, with a blank screen separating each set of images. The session began with two training items pitting an animate being (an animal) against an inanimate object. In one trial, children were asked to point to the animate referent, and on the other, they were asked to point to the inanimate referent, with sides counterbalanced, thereby establishing the premise that either referent and either side of the screen was viable. The prompt made use of the neutral predicates from the neutral condition of the unfamiliar trials. Children experienced no difficulty in the training trials.

4.3. Procedure

Children were introduced to the task as a game. They were told the experimenter was going to ask them to 'find some things,' and they would have to listen carefully to see if they could find what the experimenter was talking about. After the two training trials with the familiar words (*horse*, *spoon*), the experimenter transitioned to the test session with the six target trials. The experimenter told the child that this time, they might not recognize some of the words, but they could still try their best to figure out what the experimenter was talking about.

Each trial had the same structure. During the blank screen preceding the two images, the experimenter repeated two assertions about the target image (e.g., *Look! The mot is friendly. The mot is friendly.*). The trial then preceded to the two images, and the experimenter repeated the sentence (*The mot is friendly.*), and directed the child to make a choice (e.g., *Can you find the mot? Point to the mot!*). The experimenter recorded the child's choice as it was made by circling L or R on a response sheet.

4.4. Results

Children averaged 73% correct (SD = 0.30), which was significantly different from chance (one-sample Wilcoxon signed rank test, V = 176, $p = 0.008$ (run in R with wilcox.test, mu = 0.5), and binomial probability of correct/total responses (93/128), $p < 0.0001$). 3 children displayed an average % correct less than 0.5, two had an average of 0.5, and 17 of the 22 had an average above 0.5. Of the latter, eight of the 22 children received a perfect score. We note here that a smaller sample size of 10 children yielded results no different from chance level.

4.5. Discussion

The results of Experiment 2 had raised the possibility that 36-month-olds either did not succeed because of the time pressures of the task, or that the looking time data masked their abilities, perhaps due to the fact that their selection was fleeting and did not persist beyond a brief look to the animate referent. To test these hypotheses, we adapted our task to remove the time pressure and allow children to indicate their selection by ostension rather than infer it via their eye gaze. Quite remarkably, this strategy worked, the results revealing 36-month-olds can, in fact, deduce the potential referent for a novel nominal, given modification by a known adjective that encodes animacy constraints.

We are, then, faced with two sets of intriguing contrasting results. First, through age 24 months of age, infants can rapidly recruit a known verb with animacy constraints to identify the referent for a known noun—and yet, they are unable to do so when the known lexical item is an adjective. Second, at 36 months of age, in the same incremental processing task, young children still show no detectable signature demonstrating that they can recruit this adjectival lexical semantic information—and yet, in an interactive, offline version of the task, they are able to do so with some level of proficiency (although not at ceiling). What, then, accounts for the contrast between grammatical categories (verbs vs. adjectives) and the contrast in the nature of the task? In the next section, we offer some possible answers to these questions, which in turn, raise additional questions to be addressed in future research.

5. General discussion

A central question in language acquisition concerns the informational resources that are available for mapping an interpretation onto a novel phonological form. Simply put, how does a child learn the meaning of words? A prime source of information is the semantic representation of co-occurring words, and in particular, words that carry selectional restrictions on the other lexical items with which they compose. Previous work by Ferguson, Graf, and Waxman (2014, 2018) has shown that infants as young as 19 months of age can recruit the animacy restrictions of a known verb to identify an animate referent for a novel noun that occupies its subject argument (e.g., *The larp is eating*). Equipped with these findings, we conducted a set of experiments with 24-month-olds (the oldest age they demonstrated could successfully do this with verbs) to determine whether a parallel process holds for adjectives that carry similar selectional constraints and that are presented in the same utterance-final position (e.g., *The larp is hungry*). What we found, however, was at 24 months and even at 36 months of age, young children fail to recruit the lexical semantic information encoded in adjectives in an online task, as they do with verbs. Thus, where children are called upon to rapidly recruit known words to learn a new noun, verbs are privileged over adjectives. However, when the time pressure was removed, and the task adapted to a more interactive, offline paradigm where these children are given a chance to integrate the linguistic and visual stimuli, and deploy their reasoning process, 36-month-olds met with success, recruiting the known adjective to identify the animate referent of the novel noun.

Based on the current set of findings, two main questions arise. First, what is it about the linguistic stimuli with verbs that successfully constrains the incremental search for the intended nominal referent? (Or alternatively, What is it about adjectives that makes this process difficult?) Second, why do 36-month-olds fail to recruit a known adjective to constrain the referent selection with novel nouns in our online task, but successfully do so when the task is adapted to be more interactive and exclude a time limit? Here, we lay out some possibilities to explain each contrast, addressing each puzzle in turn.

We begin with the second question: why should 36-month-olds fail to home in on a nominal referent for the novel noun when presented with known adjectives in the online task of Experiment 2, and yet demonstrate success when presented with the same adjectives and engaged in the offline version in Experiment 3? We offer the suggestion here that these three-year-olds may know enough about the animate targets in the *Unfamiliar* trials to know that they already have a label, and this lexical semantic knowledge influenced their decision process. Faced with being asked to find the referent within a matter of seconds in a passive looking task, 36-month-olds may have found themselves

juggling their lexical knowledge of familiar nominal labels (which would have tipped the scales in the direction of the unknown label being mapped to the unknown inanimate referent (Clark, 1987; Markman & Wachtel, 1988; Merriman, Bowman, & MacWhinney, 1989) with the animacy constraint of the adjectives (which should have tipped the scales in the direction of the unknown label being mapped to an animate referent that could realistically possess the property described by the adjective). As a result, neither constraint won out within the limited amount of time they were given in the online task, and perhaps they also rapidly became bored as the Test phase progressed: consequently, their eye gaze did not lock onto a referent. Suggestive evidence that lexical knowledge is a factor in this older age group comes from occasional verbal responses from the children participating in the forced choice version of this task. Occasionally, the 3-year-olds voluntarily labeled the animals at Test (e.g., "A lizard!", "A rhino!", "A 'mingo!", or even "A mo-hog!" for *hedgehog*.) However, they still appeared to be attending to animacy, since others occasionally remarked, "He doesn't look sad," or "Why is he angry?"). Thus, in the interactive, forced-choice paradigm of Experiment 3, where children had the space to reconcile these competing streams of semantic knowledge within the more relaxed constraints of this version of the task, they succeeded in allowing the animacy constraint to do its work.

Is this, then, what also happened with 24-month-olds in Experiment 1, who failed to recruit the animacy constraints of adjectives in learning novel nouns when they had successfully done so with verbs in an otherwise identical paradigm in Ferguson et al. (2018)? It is possible that this kind of competition among constraints may also be at play in this younger group. At the same time, there may be other reasons why these younger children succeeded with verbs and not with adjectives, which we outline below. We turn now, then, to what distinguishes success with verbs from failure with adjectives.

To explain the contrast between verbs and adjectives, we entertain two hypotheses. The first posits that in these tasks, verbs may have been inherently privileged over adjectives, by virtue of their informative morphology, which provided a surface-level cue about semantic meaning. While the verbs and the adjectives both appeared in the same utterance-final position following *is* in our linguistic stimuli, they differed with respect to their morphology, which is tied to both the grammatical category and the syntax. In the previous verb experiments, children were presented with sentences such as, *The larp is eating* or *crying*. In this sentence, *is* is an auxiliary verb heading a VP, where the main verb features inflectional morphology indicating present progressive aspect (*eating*). By contrast, in the current adjective experiments, children heard sentences such as, *The larp is hungry* or *sad*. In this case, *is* is a copular verb followed by a predicate, which (at least in this instance, in this language) does not host any such morphology. This difference in the stimuli may carry an important consequence.

The verb version signals to the listener that a durative event is currently unfolding. Thus, the morphosyntactic information in the verb condition, presented first in the Dialogue phase, may be a call to attention to inspect the visual stimuli in the Test phase for an event participant (the larp) who could be the agent of this ongoing event. (See Syrett et al. (2014) and Arunachalam et al. (2016) for related discussion concerning the role of adverbial VP modifiers.) In this way, the surface-level morphosyntactic cues may serve as a bootstrap into semantic meaning—a finding consistent with a wide range of syntactic bootstrapping since Brown (1957). (See also, Fisher (1996), Gleitman (1990), Landau and Gleitman (1985), Lukyanenko and Fisher (2014); Naigles (1990).)

We might then hypothesize that if the aspectual morphology on the verb that was featured in the previous experiments is responsible (in part) for the results, then removing this surface-level cue that an event is in progress should either result in failure (as with adjectives) or the referent selection being delayed until later in the Test phase. This situation might arise in English with another tense (e.g., *The larp [ate/will eat] a carrot.*). Or, it might arise in a language other than English, which is morphosyntactically impoverished. One such example is Mandarin Chinese. By contrast, if we think that it is the lexical semantics of the verb that matters, or that verbs are inherently privileged by virtue of their argument structure (and the semantic role(s) it signals), then the same pattern of success should emerge in both of these cases.

At the same time, it is not the case that adjectives *never* display morphosyntactic cues, either cross-linguistically, or even in English. While adjectives may not host the same sort of aspectual morphology as verbs do, they do host other morphological information in languages outside of English, and even in English under other circumstances. If it is simply the presence of morphosyntactic cues that matters, calling the participant's attention to highlighted features of the scene beyond the predicate in question, then we might hypothesize that if the adjective were modified accordingly, 36- and perhaps even 24-month-olds might succeed in this task. Thus, if one turns to a language such as French or Spanish, in which adjectives obligatorily host morphological information indicating grammatical number and gender marking, children might successfully use these cues to identify the referent of a novel noun (maybe especially given the obligatory morphosyntactic agreement between the noun and its modifier).

Even in English, adjectives can be modified with suffixes such as *–ish*, which signals proximity or similarity to a standard of some sort. Given independent evidence that children at 36 months of age can rapidly recruit such morphosyntactic information (see e.g., van Heugten and Christophe (2015); Kouider, Halberda, Wood, and Carey (2006); Lew-Williams and Fernald (2007); Wood, Kouider, and Carey (2009)) and that the *–ish* morpheme can be used to assign a word an adjectival meaning (Klibanoff & Waxman, 2000; Waxman & Booth, 2003), exploring this possibility seems like a promising line of future research. If infants continue to fail in this task with adjectives under these conditions, this would be consistent with the hypothesis that there is indeed something special about the kind of event-relevant (aspectual) information verbs denote. If differences surface between these types of adjectival morphemes, then these distinctions would be a robust indicator of the potency of different kinds of morphosyntactic cues in word learning and their informativity concerning category membership, beyond being directly applicable to learning the meaning of the content words with which they combine.

We might also consider the syntactic position of the adjective as being a potential cue to meaning. Previous studies have documented that adjectives placed in prenominal (or attributive) position signal contrastive meaning (Diesendruck, Hall, & Graham, 2006; Prasada, 1992). Having presented a baseline directly comparing verbs and adjectives in utterance-final position, we might now test what participants do when they are told, *Look at the hungry larp!*

The second hypothesis regarding the difference between verbs and adjectives concerns the truth conditions of the proposition expressed by the utterance and how they align with the visual scene (i.e., the discourse context). Specifically, we might ask what kind of information presented in the visual scene at hand would allow experimental participants to evaluate the assertion expressed by the speaker's utterance as true or not. The key difference between verbs and adjectives may be one that relates to the difference between inspection of events (and event participants) and states. In the verbal version of the task, the child must inspect the scene for perceptual correlates signaling that an event is taking place, and that there is an agent engaged in that event (who in turn displays perceptual attributes of being a participant in that event). Thus, when one hears a speaker say, *The larp is eating*, the way in which to make this determination is to inspect the context at hand: is the larp (or something agentive that could be referred to as *the larp*) *eating*?

In the adjectival version (e.g., *the larp is hungry*), there is not necessarily a perceptual cue that could be a reliable indicator about the mental or psychological state of the agent. One can be hungry, sad, or tired without a stable visible manifestation of this state, and whatever manifestations there are could vary across kinds. For example, one could be smiling or staring straight ahead and be hungry or sad, and

what it looks like for a human or a hedgehog to be hungry or sad indeed may not be the same. However, there is a fairly reliable probability that being sad or angry or tired will manifest itself in a particular way; perhaps younger infants do not yet know this, while children 24 months and older do, and so these children are somewhat confused at test when they are presented with an animate being that does not exhibit those physical correlates.

This difference gives rise to the following question: if instead of verbs like *eat* or *cry*, 36-month-old children were presented with propositional attitude verbs like *think*, *want*, or *believe*, which not only lack a physical manifestation, but also pose a challenge for children through age 4–5 (Hacquard & Lidz, 2018; Harrigan, Hacquard, & Lidz, 2018; Papafragou, Cassidy, & Gleitman, 2007), would they *not* achieve the same rate of success as the younger age group, especially given the challenges faced by children at this age in tasks that tap into theory of mind (see de Villiers, 2005; Wimmer & Perner, 1983), or would they succeed, since there would be no possible mismatch between visible properties of the scene and the verb meaning? Exploring the range of known verbs for which this particular paradigm works or does not has the potential to uncover important details about the representation of lexical semantic knowledge in the child's developing lexicon and the influence of this knowledge as the child acquires new words.

A future goal for us will thus be to experimentally investigate these various hypotheses both in English and cross-linguistically in order to better understand their role in word learning. The modest yet significant findings here, set against the backdrop of the previous work, thus raise key questions about the nature of the semantic representations children have assigned to words that they comprehend and produce, and the expectations about language and the world that arise as a result of these representations. Just as importantly, this new work also paves the way for a richer discussion about what specific aspects of these words—whether it be their lexical semantics, their morphosyntactic marking and what it signals, or their truth conditional contribution—guide the deductive processes in which the learner is continuously engaged.

Acknowledgements

We gratefully acknowledge funding from NIH R01HD083310 to S. Waxman, research funding from Rutgers University to K. Syrett, comments from participants of the 2017 Rutgers Workshop on Word Learning and the audience at BUCLD 43, and the helpful feedback from three anonymous reviewers.

Appendix A

Instructions in adult 'animacy' norming task

Some things in the world are **animate**. These are things (like humans and animals) that can move on their own, can have emotions, and may be able to communicate. Some things in the world are **inanimate**. These are things that cannot move on their own, do not have emotions, and cannot communicate. Some of these things (like books, clothing, or toys) are manmade. Others (like trees, water, or stone) come from nature. Here, we want you to think about [animate/inanimate] things. Please rate each of the adjectives in the following list from 1 to 5 using the following scale:

1: cannot be used to describe [animate/inanimate] things
3: can maybe be used to describe [animate/inanimate] things
5: is ok being used to describe [animate/inanimate] things

Sample dialogues featured in experiments 1 and 2

Familiar Trials

Animate targets: bird, cow, dog; predicates: so close, very nice, so little

Inanimate targets: bottle, car, spoon; predicates: right here, nearby, very clean

(1) Target: bird

Look here!
The bird is so close. Really, the bird is so close? Yes. Let's find the bird.
Look at the bird!

Unfamiliar trials {Informative/Uninformative}

(2) Target: lizard

Look at this!
The blick is {happy/so little}. Yeah, the blick is {happy/so little}. Yes. Let's find the blick!
Look at the blick!

(3) Target: rhinoceros

Look here!
The dax is {sleepy/very nice}. Oh yes, the dax is {sleepy/very nice}. Yes. Let's find the dax!
Look at the dax!

(4) Target: flamingo

Look here!
The fisk is {sad/so close}. Really, the fisk is {sad/so close}? Yes. Let's find the fisk!
Look at the fisk!

(5) Target: hedgehog

Look here!
The larp is {hungry/right here}. Yeah, the larp is {hungry/right here}. Yes. Let's find the larp!
Look at the larp!

(6) Target: armadillo

Look at this!
The mot is {friendly/very clean}. Oh yes, the mot is {friendly/very clean}. Yes. Let's find the mot!
Look at the mot!

(7) Target: lemur

Look here!
The vep is {angry/nearby}. Really, the vep is {angry/nearby}? Yes. Let's find the vep!
Look at the vep!

Appendix B. Supplementary material

Supplementary data to this article can be found online at https://doi.org/10.1016/j.cognition.2019.104033.

References

Altmann, Gerry T. M., & Kamide, Yuki (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition, 73*(3), 247–264.
Arunachalam, Sudha, Syrett, Kristen, & Chen, Yonxiang (2016). Lexical disambiguation in verb learning: Evidence from the intransitive frame in English and Mandarin Chinese. *Frontiers in Psychology (Language Sciences), 7*, 1–14.
Arunachalam, Sudha, & Waxman, Sandra R. (2010). Meaning from syntax: Evidence from

2-year-olds. *Cognition, 114*(3), 442–446.
Bates, D., Kliegl, R., Vasishth, S., & Baayen, R. H. (2015). Parsimonious mixed models. Available from arXiv:1506.04967 [stat.ME].
Boland, Julie E. (2005). Visual arguments. *Cognition, 95*, 237–274.
Borovsky, Arielle, & Creel, Sarah C. (2014). Children and adults integrate talker and verb information in online processing. *Developmental Psychology, 50*(5), 1600–1613.
Borovsky, Arielle, Elman, Jeffrey L., & Fernald, Anne (2012). Knowing a lot for one's age: Vocabulary skill and not age is associated with anticipatory incremental sentence interpretation in children and adults. *Journal of Experimental Child Psychology, 112*(4), 417–436.
Brown, Roger W. (1957). Linguistic determinism and the part of speech. *The Journal of Abnormal and Social Psychology, 55*(1), 1–5.
Chomsky, Noam (1965). *Aspects of the theory of syntax.* Cambridge: MIT Press.
Clark, Eve V. (1987). The principle of contrast: A constraint on language acquisition. In B. MacWhinney (Ed.). *Mechanisms of language acquisition* (pp. 1–34). Hillsdale, NJ: Lawrence Erlbaum Associates.
Dale, Philip S., & Fenson, Larry (1996). Lexical development norms for young children. *Behavioral Research Methods, Instruments, & Computers, 28*(1), 125–127.
de Villiers, Jill G. (2005). Can language acquisition give children a point of view? In J. Astington, & J. Baird (Eds.). *Why language matters for theory of mind* (pp. 186–219). New York, NY: Oxford University Press.
Diesendruck, Gil, Hall, D. Geoffrey, & Graham, Susan (2006). Chidlren's use of syntactic and pragmatic knowledge in the interpretation of novel adjectives. *Child Development, 77*(1), 16–30.
Dink, Jacob W., & Ferguson, Brock (2015). eyetrackingR: An R package for eye-tracking data analysis. Retrieved from http://www.eyetracking-r.com.
Eberhard, Kathleen, Spivey-Knowlton, Michael J., Sedivy, Julie C., & Tanenhaus, Michael K. (1995). Eye movements as a window into spoken language comprehension in natural contexts. *Journal of Psycholinguistic Research, 24*(6), 409–436.
Fenson, L., Pethik, 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*, 95–116.
Ferguson, Brock, Graf, Eileen, & Waxman, Sandra R. (2014). Infants use known verbs to learn novel nouns: Evidence from 15- and 19-month-olds. *Cognition, 131*(1), 139–146.
Ferguson, Brock, Graf, Eileen, & Waxman, Sandra R. (2018). When veps cry: Two-year-olds efficiently learn novel words from linguistic contexts alone. *Language Learning and Development, 14*(1), 1–12.
Fernald, Anne, Thorpe, Kirsten, & Marchman, Virginia A. (2010). Blue car, red car: Developing efficiency in online interpretation of adjective–noun phrases. *Cognitive Psychology, 60*(3), 190–217.
Fernald, Anne, Zangl, Renate, Portillo, Ana L., & Marchman, Virginia A. (2008). Looking while listening: Using eye movements to monitor spoken language. In Irena A. Sekerina, Eva M. Fernández, & Harald Clahsen (Eds.). *Developmental psycholinguistics: On-line methods in children's language processing* (pp. 97–135). Amsterdam: John Benjamins.
Fisher, Cynthia (1996). Structural limits on verb mapping: The role of analogy in children's interpretations of sentences. *Cognitive Psychology, 31*(12), 41–81.
Gleitman, Lila R. (1990). The structural sources of verb meaning. *Language Acquisition, 1*(1), 3–55.
Goodman, Judith C., McDonough, Laraine, & Brown, Natasha B. (2008). The role of semantic context and memory in the acquisition of novel nouns. *Child Development, 69*(5), 1330–1344.
Grimshaw, Jane (1979). Complement selection and the lexicon. *Linguistic Inquiry, 10*(2), 279–326.
Grimshaw, Jane (1990). *Argument structure.* Cambridge, Massachusetts: MIT Press.
Hacquard, Valentine, & Lidz, Jeffrey (2018). Children's Attitude Problems: Bootstrapping verb meaning from syntax and pragmatics. *Mind and Language, 33*(3), 1–18.
Harrigan, Kaitlyn, Hacquard, Valentine, & Lidz, Jeffrey (2018). Three-year-olds' understanding of desire reports is robust to conflict. *Frontiers in Psychology, 9*, 119.
van Heugten, Marieke, & Christophe, Anne (2015). Infants' acquisition of grammatical gender dependencies. *Infancy, 20*(6), 675–683.
Jackendoff, Ray S. (1990). *Semantic structures.* Cambridge: MIT Press.
Jackendoff, Ray S. (1997). Twistin' the night away. *Language, 73*(3), 534–559.
Kamide, Yuki, Altmann, Gerry T., & Haywood, Sarah L. (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language, 49*(1), 133–156.
Klibanoff, Waxman, & Sandra, R. (2000). Basic level object categories support the acquisition of novel adjectives: Evidence from preschool-aged children. *Child Development, 71*(3), 649–659.
Kouider, Sid, Halberda, Justin, Wood, Justin, & Carey, Susan (2006). Acquisition of English number marking: The singular-plural distinction. *Language Learning and Development, 2*(1), 1–25.
Kuznetsova, Alexandra, Brockhoff, Per B., Christensen, Rune, & Haubo Bojesan (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software, 82*(13), 1–26.
Landau, Barbara, & Gleitman, Lila R. (1985). *Language and experience: Evidence from the blind child.* Cambridge: Harvard University Press.
Levin, Beth, & Rappaport Hovav, Malka (1991). Wiping the slate clean: A lexical semantic exploration. *Cognition, 41*(1–3), 123–151.
Lew-Williams, Casey, & Fernald, Anne (2007). Young children learning Spanish make rapid use of grammatical gender in spoken word recognition. *Psychological Science, 18*(3), 193–198.
Lukyanenko, Cynthia, & Fisher, Cynthia (2014). 30-month-olds use verb agreement features in online sentence processing. In W. Orman, & M. J. Valleau (Eds.). *Proceedings of the 38th annual Boston university conference on language development* (pp. 292–305). Somerville, MA: Cascadilla Press.
Mani, Nivedita, & Huettig, Falk (2012). Prediction during language processing is a piece of cake—But only for skilled producers. *Journal of Experimental Psychology: Human Perception and Performance, 38*(4), 843–847.
Markman, Ellen, & Wachtel, Gwyn F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology, 20*(2), 121–157.
Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods, 164*(1), 177–190.
Merriman, William E., Bowman, Laura L., & MacWhinney, Brian (1989). The mutual exclusivity bias in children's word learning. *Monographs of the Society for Research in Child Development* i-129.
Naigles, Letitia (1990). Children use syntax to learn verb meaning. *Journal of Child Language, 17*(2), 357–374.
Ninio, A. (2004). Young children's difficulty with adjectives modifying nouns. *Journal of Child Language, 31*(2), 255–285.
Papafragou, Anna, Cassidy, Kimberly, & Gleitman, Lila R. (2007). When we think about thinking: The acquisition of belief verbs. *Cognition, 105*(1), 125–165.
Pinker, Steven (1994). How could a child use verb syntax to learn verb semantics? *Lingua, 92*(1), 377–410.
Prasada, Sandeep. (1992). Acquisition of adjective meanings: A lexical semantic approach. Doctoral dissertation, MIT.
Pustejovsky, James (1995). *The generative lexicon.* Cambridge: MIT Press.
R Development Core Team (2012). *R: A language and environment for statistical computing.* Vienna, AU: R Foundation for Statistical Computing Retrieved from http://www.r-project.org.
Rappaport Hovav, Malka, & Levin, Beth (2001). An event structure account of English resultatives. *Language, 77*(4), 766–797.
Resnik, Philip (1996). Selectional constraints: An information-theoretic model and its computational realization. *Cognition, 61*(1), 127–159.
Sandhofer, Catherine M., & Smith, Linda B. (1999). Learning color words involves learning a system of mappings. *Developmental Psychology, 35*(3), 668–679.
Sedivy, Julie C., Tanenhaus, Michael K., Chambers, Craig G., & Carlson, Gregory N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition, 71*, 109–147.
Soja, Nancy N. (1994). Young children's concept of color and its relation to the acquisition of color words. *Child Development, 65*(3), 918–937.
Syrett, Kristen. (2007). Learning about the structure of scales: Adverbial modification and the acquisition of the semantics of gradable adjectives. Doctoral dissertation, Northwestern University.
Syrett, Kristen, Arunachalam, Sudha, & Waxman, Sandra R. (2014). Slowly but surely: Adverbs support verb learning in 2-year-olds. *Language Learning and Development, 10*(3), 263–278.
Syrett, Kristen, & Lidz, Jeffrey (2010). 30-month-olds use the distribution and meaning of adverbs to interpret novel adjectives. *Language Learning and Development, 6*(4), 258–282.
Tanenhaus, Michael K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, Julie C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science, 268*(5217), 1632–1634.
Thorpe, Kirsten, Baumgartner, Heidi, & Fernald, Anne (2006). Children's developing ability to interpret adjective-noun combinations. In David Bamman, Tatiana Magnitskaia, & Colleen Zaller (Eds.). *Proceedings of the 30th annual Boston-university conference on language development* (pp. 631–642). Somerville, MA: Cascadilla Press.
Tomasello, Michael (2001). Perceiving intentions and learning words in the second year of life. *Language Acquisition and Conceptual Development, 3*, 132–158.
Valian, Virginia, Prasada, Sandeep, & Scarpa, Jodi (2006). Direct object predictability: Effects on young children's imitation of sentences. *Journal of Child Language, 33*, 247–269.
Van Valin, Robert D., Jr. (1990). Semantic parameters of split intransitivity. *Language, 66*(2), 221–260.
Wagner, Katie, Jergens, Jill, & Barner, David (2018). Partial color word comprehension precedes production. *Language Learning and Development, 14*(4), 241–261.
Waxman, Sandra R., & Booth, Amy (2003). Mapping words to the world in infancy: Infants' expectations for count nouns and adjectives. *Journal of Cognition and Development, 4*(3), 357–381.
Wimmer, Heinz, & Perner, Josef (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition, 13*, 103–128.
Wood, Justin, Kouider, Sid, & Carey, Susan (2009). Acquisition of singular-plural morphology. *Developmental Psychology, 45*(1), 202–206.
Yuan, Sylvia, & Fisher, Cynthia (2009). "Really? He blicked the cat?": Two-year-olds learn distributional facts about verbs in the absence of a referential context. *Psychological Science, 20*(5), 619–626.
Yuan, Sylvia, Fisher, Cynthia, Kandhadai, Padmapriya, & Fernald, Anne (2011). You can stipe the pig and nerk the fork: Learning to use verbs to predict nouns. In Nick Danis, Kate Mesh, & Hyunsuk Sung (Eds.). *Proceedings of the 35th annual Boston university conference on language development* (pp. 665–677). Somerville, MA: Cascadilla Press.
Zucchi, Alessandro (2013). *The language of propositions and events. Issues in the syntax and semantics of nominalization.* Dordrecht: Kluwer.