THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Top-down influence in young children's linguistic ambiguity resolution.

Citation for published version:

Rabagliati, H, Pylkkänen, L & Marcus, GF 2013, 'Top-down influence in young children's linguistic ambiguity resolution.', *Developmental Psychology*, vol. 49, no. 6, pp. 1076-1089. <https://doi.org/10.1037/a0026918>

Digital Object Identifier (DOI):

[10.1037/a0026918](https://doi.org/10.1037/a0026918)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Developmental Psychology

Publisher Rights Statement:

© Rabagliati, H., Pylkkänen, L. & Marcus, G. F. Jun 2013, "Top-down influence in young children's linguistic ambiguity resolution", in *Developmental Psychology*. 49, 6, p. 1076-1089. DOI: 10.1037/a0026918, final publication available online at <http://psycnet.apa.org/journals/dev/>. This article may not exactly replicate the final version published in the APA journal. It is not the copy of record.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

RUNNING HEAD: LEXICAL RESOLUTION

Top-down influence in young children's linguistic ambiguity resolution

Hugh Rabagliati (1,2), Liina Pykkänen (1,3) & Gary F. Marcus (1)

(1) Department of Psychology, New York University

(2) Department of Cognitive, Linguistic and Psychological Sciences, Brown
University

(3) Department of Linguistics, New York University

Author Note

Hugh Rabagliati, New York University and Brown University; Liina Pykkänen, New York University; Gary F. Marcus, New York University.

Hugh Rabagliati is now at the Department of Psychology, Harvard University.

This work was supported by NIH grant R01-HD48733. Thanks to Gregory Murphy, Athena Vouloumanos, Marjorie Rhodes, Sandeep Prasada and Anja Sautmann for helpful discussion and support, and Amanda Pogue and Christina Starmans for help with testing.

Please address correspondence to Hugh Rabagliati, Department of Psychology, Harvard University, Email: hugh@wjh.harvard.edu

Abstract

Language is rife with ambiguity. Do children and adults meet this challenge in similar ways? Recent work suggests that, while adults resolve syntactic ambiguities by integrating a variety of cues, children are less sensitive to top-down evidence. We test whether this top-down insensitivity is specific to syntax, or a general feature of children's linguistic ambiguity resolution, by evaluating whether children rely largely or completely on lexical associations to resolve lexical ambiguities (e.g., the word *swing* primes the baseball meaning of *bat*), or additionally integrate top-down global plausibility. Using a picture choice task, we compared 4-year-olds' ability to resolve polysemes and homophones with a Bayesian algorithm reliant purely on lexical associations, and found that the algorithm's power to predict children's choices was limited. A second experiment confirmed that children override associations and integrate top-down plausibility. We discuss this with regard to models of psycholinguistic development.

144 words

Keywords: Language development; child language processing; lexical ambiguity; sense resolution

Top-down Influence in Young Children's Linguistic Ambiguity Resolution

Human language is rife with ambiguity. Sentences can be ambiguous between multiple syntactic structures, words can be ambiguous between multiple meanings or senses, and an acoustic signal can be ambiguous between multiple words. Nevertheless, adults manage to resolve the vast majority of linguistic ambiguity rapidly and accurately.

Research in psycholinguistics suggests that this accuracy is achieved through a cognitive architecture in which a wide variety of information sources, both bottom-up and top-down, are used to select the best interpretation from the multiple possible alternatives (see, e.g., Altmann, 1998; Dahan & Magnuson, 2006; MacDonald, Pearlmutter, & Seidenberg, 1994; Swinney, 1979). For example, a noun/verb homophone like *duck* could be resolved using information from the same level of representation, such as priming by a related word (e.g., *quack*), or through top-down feedback from its syntactic context (e.g., whether it is preceded by a definite article).

Far less is known, however, about how the mechanisms of ambiguity resolution develop. To what extent does the cognitive machinery that children use for language comprehension relate to the machinery used by adults? In fact, recent work suggests that, in many respects, child and adult language processing are similar. Like adults, children process language quickly and incrementally: They make guesses about plausible interpretations based on the currently available evidence, and do not hold off interpreting a sentence until it has finished (Fernald, Zangl, Portillo, & Marchman, 2008; Snedeker, 2009; Trueswell & Gleitman, 2004; Trueswell, Sekerina, Hill, & Logrip, 1999). But in other respects children's ambiguity resolution is very different. In particular, their

coordination of different levels of linguistic representation to ultimately choose an interpretation is distinctly non-adult-like.

Research on syntactic ambiguity resolution suggests that while children have little difficulty using bottom-up information in their decisions, they find the integration of top-down information considerably more demanding. Evidence for this distinction comes from children's ability to resolve phrases such as *tickle the frog with the feather*. Here, the ambiguity arises over where to attach the prepositional phrase *with the feather*, as an instrument phrase attached to the verb (tickling using the feather) or as a modifier of the noun (the frog holding the feather).

5-year-olds can resolve these ambiguities using bottom-up cues such as lexical statistics, the frequency with which a verb is found in a particular syntactic structure. Because a prepositional phrase subsequent to the verb *tickle* most frequently describes an instrument, children tend to attach all ambiguous prepositional phrases following *tickle* as instruments, while for verbs that typically do not take an instrument phrase (e.g., *choose the frog with the feather*), children take the phrase to modify *the frog* (Kidd & Bavin, 2005; Snedeker & Trueswell, 2004; Trueswell, et al., 1999). This bottom-up facility extends to additional cues, such as the prosodic rhythm and stress in a sentence (Snedeker & Yuan, 2008).

However, children's use of top-down information is very different. For instance, adults will quickly account for a cue called referential context. If *tickle the frog with the feather* is uttered in front of two frogs, one holding a feather and one empty-handed, adults immediately make the inference that the prepositional phrase modifies the noun in order to disambiguate which of the two frogs is referred to. In contrast, 5-year-old

children fail to make this inference and rely on lexical statistics. It is only by around 7 years that children appear to use a top-down cue like referential context (Snedeker & Trueswell, 2004; Trueswell, et al., 1999).

Building on this, Kidd and Bavin (2005) demonstrate that children aged 3;8 pay little attention to another top-down cue, global plausibility. Consider the phrase *chop the tree with the leaves*. Even though *chop* typically takes an instrument phrase, adults generally attach the prepositional phrase as a modifier of the noun, because *leaves* are implausible instruments. By contrast, their participants treated the phrase as an instrument 60% of the time, that is, they implausibly took it as an instruction to chop the tree using some leaves (see also Snedeker, Worek, & Shafto, 2009). Finally, Hurewitz, Brown-Schmidt, Thorpe, Gleitman, & Trueswell (2000) demonstrated that when 4.5-year-olds heard a question that (for adults) should promote modifier interpretations for subsequent ambiguities (e.g., *Which frog went to Mrs Squid's house?*), it did not in fact do so.

What might cause this disparity in children's use of bottom-up and top-down cues? Snedeker and Yuan (2008) discuss two possible interpretations. First, they suggest a "bottom-up hypothesis", in which young children's sentence processing shows a blanket insensitivity to top-down information. This could arise if top-down processing is too computationally arduous for children to deploy quickly and accurately, or if children do not know how to correctly align levels of representation.

Second, they discuss an informativity account, originally suggested by Trueswell and colleagues (Snedeker & Trueswell, 2004; Snedeker & Yuan, 2008; Trueswell & Gleitman, 2007; Trueswell, et al., 1999), who propose that children can potentially

integrate a variety of information sources, but that they instead choose to base their decisions on a smaller set of highly frequent and reliable cues. Under this informativity account, lexical statistics are used early because they are invariably present, they are informative, and they can be calculated via a relatively simple tabulation of frequencies. By contrast, referential context appears to be rare and less reliable (Trueswell & Gleitman, 2004), and so children rely on it less heavily than adults might.

At first glance, the viability of the bottom-up account might appear to be challenged by evidence that children can use top-down information when they are *learning* about the structure of their language, rather than processing it. Even 6-month-olds can use top-down lexical information to discover boundaries between words in fluent speech (Bortfeld, Morgan, Golinkoff, & Rathbun, 2005), and there are many demonstrations that children can use higher-level syntactic information to guide their hypotheses in word learning (Naigles, 1996). However, it is not obvious that the cognitive architectures underlying language acquisition and processing are identical. In fact, differences in the tasks involved give reason to believe that the two differ in a number of ways. During processing, children must select from a small number of known alternative interpretations, but during acquisition they have to create these interpretations themselves. This act of creation will most likely take place slowly, and rely on top-down information out of necessity (for instance, in learning the meaning of a word, the arbitrary relationship between a word's form and its meaning ensures that bottom-up information is not particularly informative). By contrast, language processing requires children to make a series of quick decisions, for which bottom-up information will be a useful guide.

In sum, evidence from language acquisition does not indubitably constrain theories of language processing development. The critical issue for both the informativity and bottom-up accounts is not whether children are blind to top-down information in toto, but why they seem to have particular difficulty integrating top-down information when resolving known linguistic ambiguities. The two proposals ascribe quite different reasons for why children's language processing architecture may differ from adults', and therefore make different predictions about ambiguity resolution in other linguistic domains. The bottom-up hypothesis predicts that top-down integration is generally difficult, and so children should ignore top-down information when resolving any type of linguistic ambiguity. By contrast, the informativity account predicts that those top-down cues that are not used for syntactic ambiguity resolution might still be integrated for other types of linguistic ambiguity.

Top-down and Bottom-up Processing in Lexical Ambiguity Resolution

One area where top-down processing might prove more relevant is the resolution of lexical ambiguity. Most words are ambiguous. This is most obviously seen with homophones (like *knight*, *bat* or *bark*) where the two meanings are completely unrelated but share a phonological form (e.g., as a result of contact between two languages). But comprehension also requires facing the subtler challenges of polysemes, words that are ambiguous between multiple related senses (e.g., *lined/academic paper*, *roasted/angry chicken*, or *birthday/playing card*). For instance, a *birthday card* and a *playing card* are different types of thing, but it does not seem to be an accident that they share a name (e.g., because they are made of the same material). The ambiguity in polysemy is not

only subtler than in homophony, it is also more common: Most frequent words are polysemous, many of them highly so (the Oxford English Dictionary lists 30 different senses for the word *line*¹).

Recent work suggests children are relatively flexible in assigning senses to words from an early age. By 4 years, they can switch between mass and count senses of nouns (*some paper/some papers*, Barner & Snedeker, 2005), extend words to novel lexical categories (e.g., *can you lipstick the trashcan?*, Bushnell & Maratsos, 1984; Clark, 1981) and resolve polysemes (Rabagliati, Marcus, & Pylkkänen, 2010; Srinivasan & Snedeker, 2011). They also have an implicit understanding that homophonous meanings are more arbitrary than polysemous senses: They assume that English homophones need not be homophonous when translated into another language, but that polysemes will still be polysemous (Srinivasan & Snedeker, 2011).

Figure 1 illustrates two ways in which ambiguous words could be resolved, one relying on information from within the lexical level, and one on top-down information based on the plausibility of different interpretations. The lexical-level route shows how children might be able to resolve such ambiguities without resort to top-down information. In particular, they could use associations between different lexical items: Spreading activation between the other words in a sentence or phrase (which we will call the context words) and the critical meanings or senses of the ambiguous word might result in increased activation of the correct sense/meaning (e.g., the word *swing* might prime the baseball meaning of *bat*). In addition, children could track co-occurrences between context words and meanings or senses, and use that information (which will be correlated with priming) for disambiguation, in a manner that is similar to their tracking

of co-occurrences between syllables (Saffran, Aslin, & Newport, 1996), words (Bannard & Matthews, 2008), and words and syntactic structures (Snedeker & Trueswell, 2004). Children's sensitivity to semantic priming (Petrey, 1977) and co-occurrence statistics (Saffran et al., 1996) suggests that neither task should be beyond them.

But there is reason to suspect that lexical associations may not be sufficient for accurate lexical ambiguity resolution. Discrimination based on lexical associations will only work well when the two meanings or senses of a word usually occur with *different* context words. This will happen if the ambiguous meanings or senses are very different, as with *bat* (Miller & Charles, 1991). But because most ambiguous words are polysemous, they have similar or related senses (e.g., the word *line* has a queue sense and an elongated-mark-on-a-page sense), and so they occur with similar context words (e.g., the phrase *the long line* is globally ambiguous). As a result, attending solely to lexical associations will often lead to difficulty determining the correct sense. To accurately resolve these senses, children need the ability to construct the potential interpretations of each sentence and determine which is more plausible. This is the second, top-down path shown in Figure 1.

The role of context in children's word recognition and processing has been quite heavily investigated, although this work has focused on the processing skills of older children learning to read, rather than younger children learning to parse spoken sentences (Simpson & Foster, 1986; Simpson, Lorschach, & Whitehouse, 1983; Stanovich, 1980; Stanovich, Nathan, West, & Vala-Rossi, 1985). Within this literature, studies focusing on lexical ambiguity have tended to find that young readers are relatively insensitive to context (e.g., Simpson & Foster, 1986). For instance, Booth, Harasaki and Burman

(2006) used a priming task in which children read aloud sentences that biased a sentence-final homophone towards one of its meanings, and then read aloud a target word that was related to one of those meanings. For 9- and 10-year-olds, reading-time for the target was unaffected by the prior context, and only 12-year-old children showed evidence that they could integrate context during lexical ambiguity resolution. But this insensitivity to context is most likely due to reading difficulties, not language processing difficulties. As evidence for this, Khanna and Boland (2010) had 7- to 10-year-old children listen to (rather than read) a prime, and then read a target out loud. When the primes were two-word phrases containing ambiguous words (e.g., *laser tag*), the reaction times of every age group varied based on whether the target was related to the selected meaning.

This suggests that, during spoken language processing, 7-year-old children can use context to resolve ambiguous words, although whether they use lexical associations or additionally integrate top-down plausibility is not clear. The informativity account provides a reason for thinking that children should be able to use top-down cues. As discussed above, lexical association cues are likely to be an unreliable guide for processing polysemous words, and so the informativity account predicts that children should turn to top-down information. And consistent with the necessity of top-down processing, there is some evidence that children are better prepared to use top-down information in resolving lexical ambiguities than syntactic ambiguities.

For example, Rabagliati et al. (2010) argued that children will sometimes assign senses to words that adults rule unlicensed, e.g., assigning a disc sense to the word *movie*. They demonstrated that 4- to 6-year-olds sometimes accept questions like *Could a movie be round?* and then explain their acceptance in terms of a shifted sense (e.g., a movie

could be round because it is on a DVD). Children presumably cannot use lexical associations or priming to assign senses that they had not previously heard. Therefore Rabagliati et al proposed that children use a process of “situational fit” to assign senses—building a representation of each potential interpretation of a phrase, and using the most plausible interpretation to assign a word sense in a top-down fashion (similar to Figure 1.)

However, the results of Rabagliati et al (2010) are not conclusive as to whether children use situational fit (and therefore top-down information) in day-to-day sense resolution. First, the sense assigned was novel, so the task was closer to word learning than ambiguity resolution. Second, only minimal association information was available for the children: The dominant sense of the word (e.g., film for *movie*) was unassociated with its context (e.g., *round*) and the to-be-shifted sense had never been encountered before. It may be that children only use top-down cues when other cues are uninformative.

This means that children’s use of top-down information in day-to-day lexical ambiguity resolution is still in question. The experiments reported here therefore test whether 4-year-old children’s ability to resolve lexical ambiguities is fully dependent on cues such as lexical associations, as the bottom-up hypothesis would predict, or whether they can go beyond this and utilize top-down information like global plausibility.

Distinctions Between Ambiguity Types

Before discussing the present experiments, there is one remaining concern. The recent studies that document children’s apparent facility in resolving the related senses of

polysemous words contrast with an earlier literature demonstrating children's difficulty resolving homophones, whose meanings are unrelated. For example, Campbell and Macdonald (1983) reported that children's accuracy resolving the subordinate (less frequent) meanings of ambiguous words was less than 20%. Beveridge and Marsh (1991) found similar results, with only a small improvement under highly constraining contexts.

Although surprising at first glance, a distinction between homophony and polysemy could conceivably be because the two types of ambiguity have been argued to be represented and accessed in different ways. In particular, while homophones are assumed to be two separate meanings that inhibit one another during lexical access, polysemous words are often argued to have a single underspecified meaning, elaborated by context (Nunberg, 1979; Rodd, Gaskell, & Marslen-Wilson, 2002). Without competitive inhibition, it should be easier to access and use less-frequent senses than (inhibited) less-frequent homophonous meanings, and this pattern has been found in adult reading time studies (Frazier & Rayner, 1990; Frisson, 2009). This difference in processing difficulty could also make polysemes comparatively easier for children, in which case our current investigation of lexical ambiguity would need to treat the two types separately.

At the same time, there are important limitations to earlier work. Campbell and Macdonald (1983) used stimuli where the subordinate meaning was very low-frequency (for example, *hair/hare*). If children did not know those meanings, accurately resolving the homophones would be extremely unlikely. Given this, we cannot say for sure whether children's differential ability is due to bona fide differences in representational format or simply due to more mundane differences in vocabulary composition. A better test would

be to directly compare the resolution of homophony and polysemy while controlling for vocabulary knowledge. The current study does exactly that.

The Current Experiments

Our first study has two parts. We assessed whether children rely on lexical associations to resolve lexical ambiguity, and also tested if resolution ability differed depending on whether the ambiguity was a homophone or polyseme. We used an offline judgment task in which 4-year-old children listened to short vignettes that served to disambiguate a target ambiguous word (e.g., *Snoopy was outside. He [chased/swung] a bat, which was big*; see Table 1). They were then shown a grid of four pictures, and were asked to choose “the picture that goes with the story.” Both the dominant (more frequent) and subordinate (less frequent) senses of the target were depicted, alongside semantic distracters for each sense.

To control for children’s vocabulary knowledge of homophones and polysemes, we excluded items that participants misidentified in a vocabulary post-test. To test whether children’s behavior could be predicted based only on the use of lexical associations, we compared their performance to a simple Bayesian algorithm (Gale, Church, & Yarowsky, 1992) which computes the probability of each sense/meaning of an ambiguous word given the other context words in the vignette, based on their previous co-occurrences within the CHILDES corpus of child language (MacWhinney, 2000). We assumed that this measure of co-occurrence statistics also provided a good proxy for conceptual associations and priming. If children attend to lexical associations over the global plausibility of a sentence, their choices should mimic the algorithm.

Experiment 2 used the same task to provide a more direct test of whether children attend to lexical associations alone, or also integrate global plausibility, by pitting the two in competition. More specifically, we compared children's accuracy at resolving ambiguous words embedded in vignettes where one sense was both lexically associated and globally plausible, compared to minimally different vignettes where one sense was lexically associated but the other was globally plausible.

Experiment 1

Experiment 1 assessed the role of lexical associations in children's lexical ambiguity resolution. To do this, we had children resolve ambiguous words embedded in vignettes, modeled what their responses should have been if they only used the lexical associations, and then compared the two.

But before assessing the role of associations, we tested whether children have more difficulty resolving the meanings of homophones than polysemes. We contrasted homophones with two types of polysemes, regular and irregular. Regular polysemes comprise a set of words whose senses fall into predictable patterns. For example, English contains (amongst others) an organism-food polysemy pattern whereby the names of plants and animals can be used to refer to the food they produce (*noisy/delicious chicken, turkey*, etc). All of our regular polysemy items were drawn from this pattern. Irregular polysemes are words whose senses are related, but do not exemplify a particular pattern, such as the senses of letter (*capital/love letter*). In our analysis, we first checked whether there was any difference in how children resolve these ambiguity types, and then assessed whether the model accurately predicted their responses.

Method

Participants.

Thirty-two 4-year-olds (range 3;10 - 4;2, $M = 4;0$, $SD = 0;1$, 16 female) were tested in a laboratory setting. An additional 9 were excluded for incorrectly answering one of the two pre-test warm-up trials. All spoke English as a first language. Children were recruited by telephone from a database of families in the New York City metropolitan area who had responded to an earlier advertisement. Not all parents provided their ethnic background; collapsing those who did in Experiments 1 and 2 ($n = 26$), our sample was 61% non-Hispanic white, with other children evenly distributed amongst other racial/ethnic groups. SES was typically mid to high.

Materials.

Two factors, sense/meaning selected by the context (dominant/subordinate) and lexical ambiguity type (homophone/irregular polyseme/regular polyseme) were varied within subjects, and one factor, the position of the disambiguating information provided by the vignette (current/prior sentence) was varied between subjects. Current-sentence context was defined as disambiguating information provided by the main verb of the sentence in which the critical noun occurred as the direct object (e.g., *Snoopy was outside. He [chased/swung] a bat, which was big*). Prior-sentence context was defined as disambiguating information provided by associated nouns and verbs in the previous sentence (e.g., *Snoopy was [reading about animals/ watching sports]. The bat was big*).

The same eight homophones, eight irregular polysemes, and eight regular polysemes (all from the organism-food pattern discussed above) were used to construct vignettes in both the current- and prior-sentence context conditions (Table 1, see Appendix 1 for a full list of the items and contexts). All of our stimuli were classified as polysemes or homophones by the Oxford English Dictionary online (<http://www.oed.com/>, see Footnote 1), except for nail (household/finger), which we treated as homophonous, rather than polysemous, because we reasoned that children would be very unlikely to perceive the relationship. We used frequency of use in CHILDES to determine which meaning/sense was dominant for each ambiguous word (see description of the method in the modeling section for procedural details). For homophones, the dominant meaning was used in 80% of cases, for irregular polysemes the dominant sense was used in 65% of cases, and for regular polysemes 75% of cases.

We produced current- and prior-sentence context vignettes for both dominant and subordinate senses. For each ambiguous word, we also created a grid of four images that depicted both the dominant and subordinate meanings/senses of the ambiguous word, along with two distracters, which were chosen because they were semantically similar to each meaning/sense, and should be known by children. Although we did not pretest whether children knew the distracters, we made sure to choose depictions of frequent words: The mean lexical frequency of the distracters (28,999 based on HAL Frequency in the English Lexicon Project, Balota et al., 2007) was 40% higher than the estimated frequency of the ambiguous word meanings (20,890). Pictures were drawn from a range of sources, were typically clip-art or illustrations, and measured approximately 20 cm².

Procedure.

At the start of each trial, the experimenter introduced children to its protagonist (*This is a story about Snoopy*). The experimenter then showed the protagonist's picture and prompted the child to name him/her. The main trial began when they completed this accurately. Children were then read the appropriate vignette, with the ambiguous word stressed (*Snoopy was outside. He chased a bat, which was big*). After the vignette, the experimenter produced the grid of four images, and asked the child to choose which one "went with" the story.

Children received 24 test trials, hearing all 24 ambiguous words, but with sense used counterbalanced between children (12 dominant and 12 subordinate per child). Trials were presented in one of two random orders, and pictures were arranged on the page in one of two random orders, making eight stimuli lists in total. Prior to test trials, participants completed two warm-up trials using unambiguous target words, and those who answered either of them incorrectly were excluded.

After the test trials, and following a 5-10 minute break, participants completed a picture-pointing vocabulary post-test on the 48 meanings tested. Children pointed to the picture that went with the word, from a selection that did not include the word's alternative sense. We excluded trials from the main analysis when participants did not know the meaning of the tested word ($\bar{M} = 4.0$ per child, $SD=2.3$), and also when participants chose a semantic distracter ($\bar{M} = 2.9$ per child, $SD=2.4$).

Results

We first tested if children have more difficulty resolving homophones than polysemes, in case we needed to account for such a difference in how we modeled their use of lexical associations. For both current- and prior-sentence contexts, we analyzed whether children appropriately changed their choice of sense/meaning depending on which was selected by the context, and whether this varied across ambiguity types. We did this using mixed-effects logistic regression models with random intercepts for subjects and items, which are more appropriate for binary data than ANOVA, and more robust to missing data. In our regressions, outcome 1 was choosing the dominant sense/meaning, and outcome 0 the subordinate sense/meaning.

For both current- and prior-sentence contexts, Figure 2 plots the proportion of time children chose the dominant sense/meaning (as opposed to the subordinate sense/meaning), split by whether the context selected for the dominant or subordinate sense/meaning. As can be seen, when context selected the dominant sense/meaning children chose it on a relatively high proportion of trials across all ambiguity types (mean proportion of dominant choices ($\underline{M}_{\text{dominant}}$) between 0.68 and 0.94). But critically, and contra Campbell and Macdonald (1983), the 4-year-old children we studied reliably changed their selection to the subordinate meaning of a homophone when the context selected it (Current: $\underline{M}_{\text{dominant}} = .25$ (0.31), $\beta = -3.3$, s.e. = 0.79, $z = 4.5$, $p < .01$; Prior: $\underline{M}_{\text{dominant}} = .44$ (0.36) $\beta = -2.0$, s.e. = 0.95, $z = 2.2$, $p = .03$), and, with one exception, this did not differ across the other ambiguity types or contexts. The sole anomaly was that children switched to the subordinate sense for the regular polysemes under current-sentence context less often than they did for the homophone condition (Current: $\underline{M}_{\text{dominant}}$

= .69 (0.30), $\beta = 2.44$, s.e. = 1.02, $z = 2.3$, $p = .02$). Post-hoc norming with adults indicated that this was because both pictures were considered relatively appropriate.²

We additionally analyzed the whole dataset in terms of accuracy, rather than sense chosen, and unsurprisingly discovered that children were reliably less accurate at choosing the subordinate meaning or sense when selected under current-sentence context, as compared to the dominant meaning or sense ($\underline{M}_{\text{dominant}} = 0.80$ (0.17), $\underline{M}_{\text{subordinate}} = 0.62$ (0.18), $\beta = -1.18$, s.e. = 0.52, $z = 2.3$, $p = .02$), although this effect was only marginal under prior-sentence context ($\underline{M}_{\text{dominant}} = 0.77$ (0.11), $\underline{M}_{\text{subordinate}} = 0.68$ (0.19), $\beta = -0.84$, s.e. = 0.53, $z = 1.60$, $p = .11$). In addition, children were more accurate than chance across all conditions (all $t > 7.0$, all $p < .01$, analyzed using one-sample t-tests against a null mean of 0.25).

In summary, children were reliably able to resolve ambiguous words, and this did not depend upon whether those words were homophones or polysemes. Next, we turn to how exactly they resolved such ambiguities: Were they primarily reliant on lexical associations, or could they integrate top-down information as well?

Children's use of lexical associations: Modeling of Experiment 1

In order to predict quantitatively how children should behave if they only use lexical associations such as priming and co-occurrences we appropriated a model from computational linguistics (Gale et al., 1992). The model uses Bayes' rule and a corpus analysis to estimate which sense/meaning an ambiguous word should take based on a) the

relative frequency of each sense/meaning, and b) the other words surrounding that ambiguous word in each vignette, and in particular, how frequently those words have previously co-occurred with each sense/meaning. This statistic therefore directly measures co-occurrences and should also be strongly correlated with the amount of priming between words (strong primes are likely to co-occur more frequently than weak primes).

We used a “bag of words” model, so named because it assumes that each context word is conditionally independent (an assumption that greatly simplifies computation, but is not entirely plausible). In short, for an ambiguous word w (e.g., *bat*) with a sense s_i (*baseball bat*) embedded in a vignette made up of a set of context words c , (e.g., *snoopy, was, outside, he, chased, which, fun*), the model estimates the probability of a particular sense given the context words, $p_w(s_i|c)$. It does this using Bayes rule, combining the prior probability of that sense of the word, $p_w(s_i)$, with the probability of the vignette’s context words appearing in a sentence or phrase containing that sense, $p_w(c|s_i)$, and the probability of the context words appearing in a sentence or phrase containing the ambiguous word independent of the sense, $p_w(c)$ (equation (2)). Because we assume that context words are conditionally independent, the conditional probability of the context words given a sense, $p_w(c|s_i)$, is calculated as the product of the probabilities of each context word given that sense (equation (1)).

$$(1) \quad p_w(c|s_i) = \prod_{j=1}^n p_w(c_j|s_i)$$

$$(2) \quad p_w(s_i|c) = \frac{p_w(c|s_i)p_w(s_i)}{p_w(c)}$$

Method.

We used the maximum likelihood method to estimate the terms of our model, based on frequencies calculated from the British and American CHILDES corpora. To calculate the prior probability of each sense/meaning, $p_w(s_i)$, we extracted every occurrence of the 24 ambiguous words, then further extracted the first occurrence from each independent conversation, and finally sense-tagged these first uses, up to up to a maximum of 75 tags per word (except for homophones with different spellings, e.g., *knight/night*, where we could use all instances). On average this resulted in 70 tags per word (SD = 13). The items *shrimp* and *herb* were excluded from further analysis as they only occurred 10 and 5 times in the corpora respectively.

Our context words, c , consisted of the open class words, pronouns and prepositions in our vignettes. To calculate the probability that a particular context word (e.g., *swing*) occurred in the presence of an ambiguous word like *bat* we summed co-occurrences in individual lines of dialogue (up to one co-occurrence per independent conversation) and divided by the overall frequency of the word. To calculate the conditional probability of a context word given a particular sense/meaning of the ambiguous word ($p_w(c_j|s_i)$), we summed co-occurrences of the context word and that particular sense/meaning, and divided by the overall frequency of the sense/meaning. For example, to calculate the probability of the word *swing* given the baseball meaning of *bat*, we counted the co-occurrences of *swing* and *bat* in its baseball meaning, and divided this by the overall frequency of the baseball meaning of *bat* (estimated as the frequency of *bat* multiplied by the prior probability of the baseball meaning).³

Results.

To what extent was children's performance reflected in the model? The bottom-up hypothesis proposes that children should pay close attention to lexical associations, in accord with their dependence on lexical associations for syntactic ambiguity resolution (Snedeker & Trueswell, 2004). And indeed, the children tested here had a very similar overall level of accuracy to the model. For the current-sentence contexts, the model's percent correct was 69%, while the children's accuracy was 70%. For the prior-sentence contexts, the model's accuracy was 67%, while children's accuracy was 68%. In addition, the model, like the children, did better when context selected for the dominant sense/meaning (Current- Model: 75%, Child: 80%; Prior- Model: 72%, Child: 78%⁴) than when context selected for the subordinate sense/meaning (Current- Model: 64%, Child: 69%; Prior- Model: 61%, Child: 57%). At a coarse level of detail, therefore, the model provided a reasonable fit to the children's data.

Next we evaluated if the model still provided a good fit at a deeper level of analysis, testing whether the two learners—model and child—made the same mistakes. We performed an item-based correlation between children's proportion of correct answers and the probability that the model assigned to the correct meaning. We found similar results in both the current and prior-sentence contexts. In each case the two patterns of answers were correlated, reliably so for the current-sentence context ($r = 0.32$, $t(42) = 2.17$, $p = .04$), and marginally for the prior-sentence context ($r = 0.29$, $t(42) = 1.94$, $p = .06$) (see Figure 3). The fact that the model had predictive capability suggests that children do use bottom-up cues. But this capability was also only moderate and

marginal, which suggests that children might be using other cues in addition, potentially top-down cues such as global plausibility.

Discussion

Experiment 1 provided two main results. First, children resolve the meanings of ambiguous words quite accurately, and this accuracy does not seem to depend on whether the ambiguous word is a polyseme or homophone. Second, children's accuracy on the different items was predicted, but only partially, by a Bayesian model using lexical associations. The model reliably predicted children's choices in the current-sentence context condition, but was only a marginally reliable guide to children's choices in the prior-sentence context condition.

Overall, the rough fit of the model to the data suggests that children do track and use associations, but can also recruit additional cues for sense/meaning resolution. In particular, children may use a top-down cue such as global plausibility, which would be contra the bottom-up hypothesis. To more directly test children's ability to use global plausibility, we conducted a second experiment that pit plausibility against lexical associations. If children can use top-down plausibility information, they should be able to override the associations.

Experiment 2

Experiment 2 compared 4-year-olds' ability to correctly resolve ambiguous words when both lexical associations and plausibility were consistent in selecting the same sense/meaning, and when they conflicted, each selecting for a different sense/meaning.

Table 2 displays sample items, in which the lexical associations between the context and the meanings of *knight/night* are extremely similar, but the plausibility of each meaning varies (e.g., Consistent: *Elmo watched a funny movie about a castle, and a princess, and a silly dragon. And there was a funny knight.* Inconsistent: *Elmo watched a funny movie about a castle, and a princess, and a silly dragon. That was a funny night*). If children resolve ambiguous words using top-down plausibility, they ought to assign reliably different senses/meanings across sentence pairs of this sort. By contrast, if children resolve ambiguous words using lexical associations, they should assign the same sense/meaning to both sentences.

We used the same picture-pointing task as before, with a mixture of 6 homophones and irregular polysemes that were collapsed in our analysis. In addition we compared children's performance with both the association-based algorithm used in Experiment 1 and adult judgments, in order to confirm our manipulation. One concern was that the sentences were more complicated than in Experiment 1, and any difficulty that children had with the task could result from that. As a control, we therefore included a set of matched sentences where the target word was unambiguous; difficulties with the task should be reflected in both conditions.

Method

Participants.

Sixteen 4-year-olds ($M = 3;11$, $SD = 0;1$, 8 female) were tested in a laboratory, as were 19 college-age adults who received course credit for participation. An additional four children were excluded for incorrectly answering a pre-test warm-up question. All

spoke English as a first language. See Experiment 1 for details of children's demographic information and recruitment procedures.

Materials.

We constructed two minimal pairs of vignettes for each of 6 target ambiguous words (3 homophones and 3 irregular polysemes, see Table 2). We chose items based on the ease of constructing the vignettes. Lexical associations between the context words and the target's senses were approximately constant and always pointed to the subordinate meaning. For each pair, the subordinate meaning was not only more associated but also more plausible in one vignette (e.g., *Elmo watched a funny movie about a castle, and a princess, and a silly dragon. And there was a funny knight.*), while in the other vignette, the association statistics were almost identical, but the dominant meaning was more plausible (e.g., *Elmo watched a funny movie about a castle, and a princess, and a silly dragon. That was a funny night.*). The control pairs were similar, but contained unambiguous target words. That is, one control item had an associated meaning (e.g., *Big Bird was visiting a castle. He saw both a sword and a horse with the jester.*) and one item had an unassociated meaning (e.g., *Big Bird was visiting a castle. He saw both a sword and a horse during the morning.*). Control words were frequency-matched to the ambiguous senses (Ambiguous mean log-frequency = 9.4 (s.d. = 1.8), Unambiguous: $\bar{M} = 8.9$ (1.9), $t(19) = 1.5$, *ns*). We counterbalanced which pair contained ambiguous targets or unambiguous controls between participants. Using a Latin square design, each participant received 3 items from each condition (ambiguous/control crossed with associated/unassociated) making 4 lists of stimuli, each presented in a different

random order. Images were selected in the same manner as in Experiment 1, and the arrangement of pictures on the page randomly varied between two different lists, making 8 lists in total. Items are listed in Appendix 2.

Procedure.

The basic trial structure was equivalent to Experiment 1. Each child heard 6 test vignettes (3 with a plausible subordinate meaning, and 3 with a plausible dominant meaning), and 6 control vignettes, counterbalanced so that only one sentence from each pair was used per list. Children also performed 3 pre-test warm-up trials containing unambiguous targets, and were excluded if they failed any. We again analyzed the data with mixed effects logistic regressions including random intercepts for subjects and items. The small number of items prevented an analysis of homophony-polysemy differences.

Adults were tested slightly differently. In order to derive a more sensitive measure of the effects of plausibility and association in these stimuli, we had them rate the acceptability of each picture on a 1-7 scale (the two target pictures only, distracter pictures were not tested) for both the ambiguous and unambiguous control items. As our dependent measure, we calculated the rating of the dominant sense as a proportion of the summed ratings of both senses, and analyzed this using a mixed effects linear regression with random subject and item intercepts. Finally, we calculated association statistics for the algorithm in the same manner as Experiment 1.

Results

Figure 4 displays the results of Experiment 2. The adults confirmed our plausibility manipulation. In both the ambiguous and unambiguous conditions, when the subordinate sense was both associated and plausible, subjects gave it a higher rating than the dominant sense (and so the acceptability ratio was low, Ambiguous: $\underline{M} = .21$ (0.21); Unambiguous: $\underline{M} = .09$ (0.18)). But the ratio score reliably increased when the dominant sense was more plausible, even though it was statistically unassociated in the ambiguous condition (Ambiguous: $\underline{M} = .74$ (0.21), $\beta = 0.53$, s.e. = 0.04, $t = 13.6$, $p < .01$; Unambiguous: $\underline{M} = .91$ (0.13), $\beta = 0.82$, s.e. = 0.03, $t = 27.12$, $p < .01$).

As expected, the algorithm reliably assigned senses based on associations: It was correct for the unambiguous controls, but chose the statistically associated meaning for the ambiguous items, not the most plausible (see Figure 4).

Finally, we tested children's ability to use plausibility and associations. As can be seen in Figure 4, when the ambiguous subordinate sense was both associated and plausible, children chose it on a high proportion of trials ($\underline{M}_{\text{dominant}} = .21$ (95% confidence interval = .11 - .29)). However, when the dominant sense became globally more plausible, children chose it reliably more often, even though it was still lexically unassociated ($\underline{M}_{\text{dominant}} = .39$ (95% c.i. = .26 - .52), $\beta = 1.27$, s.e. = 0.62, $z = 2.07$, $p = .039$). In addition, children showed a similar pattern to adults on the unambiguous items ($\beta = 4.28$, s.e. = 0.86, $z = 4.99$, $p < .01$). This indicates that children do use global plausibility in resolving lexical ambiguities, and do not exclusively rely on lexical associations.

Nevertheless, children did also use associations: They chose the plausible but unassociated dominant sense at a lower rate (39%) than they did in Experiment 1 (80%),

suggesting that children resolve lexical ambiguities by integrating both global plausibility and associations.

General Discussion

Adults process language quickly and accurately. In particular, they rapidly resolve ambiguities by integrating both bottom-up and top-down cues. How does this expert skill develop? Studies of syntactic comprehension suggest the possibility that the development of linguistic ambiguity resolution might be characterized by a broad dependence on bottom-up cues over top-down information. To evaluate this, we tested whether that imbalance extends to children's resolution of lexical ambiguities, and instead found that even 4-year-old children used a top-down cue like global plausibility information to resolve word senses. That is not to say that children did not use bottom-up cues at all. In Experiment 1 the Bayesian model offered a limited ability to predict children's choices, and in Experiment 2 children's choices were clearly swayed by the overwhelming association information. But critically, these results are not consistent with any theory that assumes children's linguistic ambiguity resolution has a blanket bottom-up quality. This suggests that children's language processing differs from adults' by degree, not kind.

These results *are* consistent with previous work that assumed children could resolve senses and meanings on the basis of global plausibility, such as Rabagliati et al.'s (2010) proposal that children use situational fit in sense resolution, assigning senses based on their fit with a partial semantic context. They also fit within the framework of the informativity account of syntactic ambiguity resolution (Trueswell & Gleitman,

2004). Trueswell proposes that children start to apply different cues at different ages as a result of tracking each cue's general reliability. As we argued in the introduction, there are reasons to suspect that whereas cues such as lexical statistics are more informative than global plausibility for syntactic ambiguity resolution, they may be less useful for lexical ambiguity resolution. This might force children to assign higher weight to top-down plausibility information in order to resolve lexical ambiguities.

Of course, the informativity account is not the only possible explanation for the difference between children's use of top-down information in syntactic and lexical ambiguity resolution. One possibility is that skills learned during language acquisition might transfer to language processing. Lexical ambiguity resolution resembles the act of word learning quite closely, in that both require the child to use the surrounding context to select between a set of candidate meanings. The major difference, of course, is that the set of candidate meanings is much larger during word learning, but the structure of the problem seems similar enough that the skills developed in one could plausibly transfer to the other. Since many theories of word learning rely on the child being able to determine the meaning of a novel word based on its syntactic/semantic context and the scene (Gleitman, 1990; Pinker, 1994), the use of this sort of plausibility information may well transfer from word learning to lexical ambiguity resolution. By contrast, it is not so obvious that plausibility (or referential context) is particularly important for syntactic development, and as such the use of plausibility for syntactic processing may be underdeveloped as compared to its use for lexical processing.

What this means in sum is that our results cannot be explained by accounts of children's language processing that postulate a blanket top-down insensitivity (such as

the bottom-up account described by Snedeker & Yuan, 2008). Rather, they have to be explained under theories in which children are differentially sensitive to different types of cues for different types of linguistic ambiguities. Given that a variety of such theories exist, further work will be necessary to pare down the members of this set.

In their use of top-down plausibility, the children tested here generally behaved in a relatively adult-like way, but it is still notable that they had a lower accuracy resolving subordinate senses/meanings than dominant senses/meanings in Experiment 1. This suggests that children have difficulty fully integrating contextual cues, and rely too heavily on each sense/meaning's prior probability. To some extent, then, this is consistent with a weaker version of Campbell and Macdonald's (1983) claim that children cannot resolve the subordinate meanings of homophones, in which subordinate meanings are simply dispreferred, not ruled out entirely.

We see a number of potential explanations for this result. One possibility is that children may have difficulty retrieving the subordinate meanings of words from memory. Alternately, it could reflect a difficulty modulating the activity of each retrieved meaning, perhaps because of under-developed executive function abilities. The role of executive function in language-processing development is currently a topic of intense interest (Choi & Trueswell, 2010; Mazuka, Jincho, & Oishi, 2009; Novick, Trueswell, & Thompson Schill, 2010). The majority of work in this area has centered on children's failure to revise initial parses of temporally ambiguous sentences that are eventually resolved to a less-frequent interpretation (as in Trueswell et al., 1999). Novick and colleagues (Novick, et al., 2010; Novick, Trueswell, & Thompson-Schill, 2005) argue that this "kindergarten-path effect" is due to children lacking the ability to inhibit highly active representations,

and they suggest that children should exhibit similar behaviors for similar reanalysis tasks, such as resolving an ambiguous word to its less-frequent meaning. Consistent with this, Khanna and Boland (2010) show that 7- to 10-year-old children's scores on a battery of executive function tests predict their ability to resolve ambiguous words. Although not intended to test the role of executive function, our participants' greater difficulty resolving less-frequent meanings/senses is consistent with Novick and colleagues' proposal. Note, though, that these differences were much smaller than Trueswell et al.'s (1999) syntactic kindergarten-path effect. This could be because executive function plays a smaller role in lexical ambiguity resolution, or it could be because the disambiguating information in our stimuli always preceded the ambiguity, minimizing the reanalysis demands.

Beyond reanalysis, there have also been suggestions (Mazuka, et al., 2009) that children's executive function difficulties may explain their failure to integrate top-down constraints. To the best of our knowledge, there is little direct evidence that executive function plays a facilitatory role in the use of top-down, but not bottom-up, information during sentence processing, and we do not take our data to be consistent with this proposal. In fact, if executive function is domain-general (which most researchers assume it to be), it would be very surprising if its underdevelopment impaired top-down processing in syntactic ambiguity resolution, but not in lexical ambiguity resolution.

One methodological note concerns the picture selection task used in both studies. Picture selection is an offline task, which measures the factors that affect final interpretations of ambiguous words, rather than the moment-by-moment processes of ambiguity resolution. Our results suggest that children and adults use similar information

sources in resolving ambiguous words, but it is still open as to whether the two populations use the same processes to integrate that information online, and future work will address that topic. To assess whether different cues are integrated in the same manner we would need to use an online measure, like eye-tracking. It could be that all cues are integrated at the earliest moment they become available (as would be expected under a fully interactive model of sentence processing development, such as the informativity account), but there are alternatives. In particular, children's integration of top-down cues may be delayed until a second reanalysis stage (cf. two-stage parsing theories, e.g., Frazier, 1987), or our offline task may have provided children with enough reflection time to use top-down cues in the final moments of their decision. Our data, therefore, leave open a window of hope for a modified version of the bottom-up hypothesis. In this version, children's initial parsing decisions are governed by bottom-up information, and top-down cues are integrated later if they are forced to make explicit decisions. Of course, this account would still have to explain why our task allowed children to use a decision stage, while the syntactic processing task from Snedeker and Trueswell (2004) did not. One possibility is that our task presented the choice in a more explicit way.

Further work will also be needed to confirm the second main result of the present study, that children resolve homophones with the same accuracy as polysemes. While we found no gross interpretive differences, it is possible that more fine-grained online measures may be able to track such a developmental dissociation. We attribute our divergence from previous results to the efforts made here to use homophones and polysemes that children knew, rather than low frequency items. Although this null result

could be taken as support for theories in which homophony and polysemy are represented and resolved in the same way (Foraker & Murphy, 2009; Klein & Murphy, 2001; Murphy, 2007) we would caution against immediately leaping to that conclusion. In particular, we did not test the condition in which homophones and polysemes have previously diverged in adults, reanalysis contexts in which the disambiguating region falls after the critical ambiguous word.

Conclusion

This research demonstrates that children's lexical ambiguity resolution is sensitive to a variety of information sources, including lexical associations but also top-down global plausibility information. That is to say, children's lexical ambiguity resolution appears similar to adults', and this is inconsistent with theories based on syntactic processing that attribute a blanket insensitivity to top-down information. Instead, this less-categorical pattern of successes and failures suggests that children's ambiguity resolution is dependent on learning the value of individual information sources, and that the value of each information source varies for different types of linguistic ambiguity. More broadly, our data present a picture of information-processing development in which children are minimally constrained in which cues they can learn to integrate, and in fact can respond to the nature of the task when determining the most relevant cues to use.

References

- Altmann, G. (1998). Ambiguity in sentence processing. *Trends in Cognitive Sciences*, 2, 146-152. doi:10.1016/S1364-6613(98)01153-X
- Balota, D., Yap, M., Cortese, M., Hutchison, K., Kessler, B., Loftis, B. (2007). The English lexicon project. *Behavior Research Methods*, 39, 445-459. doi:10.3758/BF03193014
- Bannard, C., & Matthews, D. (2008). Stored word sequences in language learning. *Psychological Science*, 19, 241-248. doi:10.1111/j.1467-9280.2008.02075.x
- Barner, D., & Snedeker, J. (2005). Quantity judgments and individuation: evidence that mass nouns count. *Cognition*, 97, 41-66. doi:10.1016/j.cognition.2004.06.009
- Beveridge, M., & Marsh, L. (1991). The influence of linguistic context on young children's understanding of homophonic words. *Journal of Child Language*, 18, 459-467. doi:10.1017/S0305000900011168
- Booth, J. R., Harasaki, Y., & Burman, D. D. (2006). Development of lexical and sentence level context effects for dominant and subordinate word meanings of homonyms. *Journal of Psycholinguistic Research*, 35, 531-554. doi:10.1007/s10936-006-9028-5
- Bortfeld, H., Morgan, J., Golinkoff, R., & Rathbun, K. (2005). Mommy and me: Familiar names help launch babies into speech stream segmentation. *Psychological Science*, 16, 298-304. doi: 10.1111/j.0956-7976.2005.01531.x

- Bushnell, E. W., & Maratsos, M. P. (1984). "Spoonng" and "Basketng": Children's dealing with accidental gaps in the lexicon. *Child Development*, 55, 893-902.
doi:10.2307/1130140
- Campbell, R. N., & Macdonald, T. (1983). Text and context in early language comprehension. In M. C. Donaldson, R. Grieve & C. Pratt (Eds.), *Early childhood development and education* (pp. 115-126). Oxford, UK: Blackwell Books.
- Choi, Y., & Trueswell, J. C. (2010). Children's (in) ability to recover from garden paths in a verb-final language: Evidence for developing control in sentence processing. *Journal of Experimental Child Psychology*, 106, 41-61.
doi:10.1016/j.jecp.2010.01.003
- Clark, E. V. (1981). Lexical innovations: How children learn to create new words. In W. Deutsch (Ed.), *The child's construction of language* (pp. 299-328). London: Academic Press.
- Dahan, D., & Magnuson, J. S. (2006). Spoken word recognition. In M. J. Traxler & M. A. Gernsbacher (Eds.), *Handbook of Psycholinguistics* (2nd ed., pp. 249-283). New York, NY: Elsevier.
- 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). Philadelphia, PA: John Benjamins.

- Foraker, S., & Murphy, G. L. (2009). *Polysemy in sentence comprehension: Effects of meaning dominance*. Paper presented at the 22nd Annual CUNY Conference on Sentence Processing, Davis, CA.
- Frazier, L. (1987). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and Performance XII: The psychology of reading* (pp. 554-586). Hillsdale, NJ: LEA.
- Gale, W., Church, K., & Yarowsky, D. (1992). A method for disambiguating word senses in a large corpus. *Computers and the Humanities*, 26, 415-439.
doi:10.1007/BF00136984
- Gleitman, L. (1990). The Structural Sources of Verb Meanings. *Language Acquisition*, 1, 3-55. doi:10.1207/s15327817la0101_2
- Hurewitz, F., Brown-Schmidt, S., Thorpe, K., Gleitman, L. R., & Trueswell, J. C. (2000). One frog, two frog, red frog, blue frog: Factors affecting children's syntactic choices in production and comprehension. *Journal of Psycholinguistic Research*, 29, 597-626.
- Khanna, M. M., & Boland, J. E. (2010). Children's use of language context in lexical ambiguity resolution. *The Quarterly Journal of Experimental Psychology*, 63, 160-193.
- Kidd, E., & Bavin, E. L. (2005). Lexical and referential cues to sentence interpretation: an investigation of children's interpretations of ambiguous sentences. *Journal of Child Language*, 32, 855-876. doi:10.1017/S0305000905007051
- Klein, D. E., & Murphy, G. L. (2001). The representation of polysemous words. *Journal of Memory and Language*, 45, 259-282. doi:10.1006/jmla.2001.2779

- MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, 101, 676-703.
doi:10.1037//0033-295X.101.4.676
- MacWhinney, B. (2000). *The CHILDES project: Tools for analyzing talk*. Hillsdale, NJ: LEA.
- Mazuka, R., Jincho, N., & Oishi, H. (2009). Development of executive control and language processing. *Language and Linguistics Compass*, 3, 59-89.
doi:10.1111/j.1749-818X.2008.00102.x
- Miller, G. A., & Charles, W. G. (1991). Contextual correlates of semantic similarity. *Language and Cognitive Processes*, 6, 1-28. doi:10.1080/01690969108406936
- Murphy, G. L. (2007). Parsimony and the psychological representation of polysemous words. In M. Rakova, G. Petho & C. Rákosi (Eds.), *The cognitive basis of polysemy*. (pp. 47-70). Fankfurt am main, Germany: Peter Lang Verlag.
- Naigles, L. (1996). The use of multiple frames in verb learning via syntactic bootstrapping. *Cognition*, 58, 221-251. doi:10.1016/0010-0277(95)00681-8
- Novick, J. M., Trueswell, J. C., & Thompson Schill, S. L. (2010). Broca's area and language processing: Evidence for the cognitive control connection. *Language and Linguistics Compass*, 4, 906-924. doi:10.1111/j.1749-818X.2010.00244.x
- Novick, J. M., Trueswell, J. C., & Thompson-Schill, S. L. (2005). Cognitive control and parsing: Reexamining the role of Broca's area in sentence comprehension. *Cognitive, Affective, & Behavioral Neuroscience*, 5, 263-281.
doi:10.3758/CABN.5.3.263

- Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, 3, 143-184. doi:10.1007/BF00126509
- Petrey, S. (1977). Word associations and the development of lexical memory. *Cognition*, 5, 57-71. doi:10.1016/0010-0277(77)90017-8
- Pinker, S. (1994). How could a child use verb syntax to learn verb semantics? *Lingua*, 92, 377-410. doi:10.1016/0024-3841(94)90347-6
- Rabagliati, H., Marcus, G. F., & Pykkänen, L. (2010). Shifting senses in lexical semantic development. *Cognition*, 117, 17-37. doi:10.1016/j.cognition.2010.06.007
- Rodd, J., Gaskell, G., & Marslen-Wilson, W. (2002). Making sense of semantic ambiguity: Semantic competition in lexical access. *Journal of Memory and Language*, 46, 245-266. doi:10.1006/jmla.2001.2810
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926-1928. doi:10.1126/science.274.5294.1926
- Simpson, G. B., & Foster, M. R. (1986). Lexical ambiguity and children's word recognition. *Developmental Psychology*, 22, 147. doi:10.1037//0012-1649.22.2.147
- Simpson, G. B., Lorchbach, T. C., & Whitehouse, D. (1983). Encoding and contextual components of word recognition in good and poor readers. *Journal of Experimental Child Psychology*, 35, 161-171. doi:10.1016/0022-0965(83)90076-0
- Snedeker, J. (2009). Sentence processing. In E. Bavin (Ed.), *The Cambridge handbook of child language* (pp. 321-338). Cambridge, UK: CUP.
- Snedeker, J., & Trueswell, J. C. (2004). The developing constraints on parsing decisions: The role of lexical-biases and referential scenes in child and adult sentence

- processing. *Cognitive Psychology*, 49, 238-299. doi: DOI
10.1016/j.cogpsych.2004.03.001
- Snedeker, J., Worek, A., & Shafto, C. (2009). *The role of lexical bias and global plausibility in children's online parsing: A developmental shift from bottom-up to top-down cues*. Paper presented at the 34th Boston University Conference on Language Development Boston, MA.
- Snedeker, J., & Yuan, S. (2008). Effects of prosodic and lexical constraints on parsing in young children (and adults). *Journal of Memory and Language*, 58, 574-608. doi: 10.1016/j.jml.2007.08.001
- Srinivasan, M., & Snedeker, J. (2011). Judging a book by its cover and its contents: The representation of polysemous and homophonous meanings in four-year-old children. *Cognitive Psychology*, 62, 245. doi:10.1016/j.cogpsych.2011.03.002
- Stanovich, K. E. (1980). Toward an interactive-compensatory model of individual differences in the development of reading fluency. *Reading Research Quarterly*, 16, 32-71. doi:10.2307/747348
- Stanovich, K. E., Nathan, R. G., West, R. F., & Vala-Rossi, M. (1985). Children's word recognition in context: Spreading activation, expectancy, and modularity. *Child Development*, 1418-1428. doi:10.2307/1130461
- Swinney, D. A. (1979). Lexical access during sentence comprehension:(Re) consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18, 645-659. doi:10.1016/S0022-5371(79)90355-4
- Trueswell, J. C., & Gleitman, L. (2004). Children's eye movements during listening: Developmental evidence for a constraint-based theory of sentence processing. In

- J. M. Henderson & F. Ferreira (Eds.), *The interface of language, vision, and action: Eye movements and the visual world* (pp. 319-346). New York, NY: Psychology Press.
- Trueswell, J. C., & Gleitman, L. R. (2007). Learning to parse and its implications for language acquisition. In M. G. Gaskell (Ed.), *Oxford handbook of psycholinguistics*. (pp. 635-657). Oxford, UK: OUP.
- Trueswell, J. C., Sekerina, I., Hill, N. M., & Logrip, M. L. (1999). The kindergarten-path effect: Studying on-line sentence processing in young children. *Cognition*, 73, 89-134. doi:10.1016/S0010-0277(99)00032-3

Footnotes

¹ The prescriptive method for distinguishing between homophones and polysemes is consulting a dictionary. Homophones are listed as separate entries, while polysemous senses are listed within an entry.

² 16 adults saw both target pictures on a computer screen alongside the relevant vignette, and moved a slider on a continuous scale between the two picture to indicate which they thought was most appropriate, with the center of the scale marked as “equally appropriate”. We analyzed the data with multiple regressions. For the current-sentence context condition, we observed an interaction such that participants’ ratings were reliably closer to the center for the regular polysemy condition, but only when the subordinate meaning was selected ($t > 2.3$, $p < .05$).

³ To avoid contexts with probability 0, we assumed that context words that did not co-occur in our corpus actually co-occurred with a frequency of 0.01 (meaning they had very low probability).

⁴ Note that children’s scores are slightly different than in our first analysis of the behavioral data, because items were excluded during modeling, and these means are item-wise, not subject-wise.

LEXICAL RESOLUTION

Table 1.

Example stimuli for each condition in Experiment 1. Dominant sense and selecting context are in bold type, while subordinate sense and context are underlined. Distracter targets are in plain type.

Current- sentence Context	Homophone	<i>Vignette</i>	Snoopy was outside. He [chased / <u>swung</u>] a <i>bat</i> , which was big.			
		<i>Targets</i>	Animal bat	<u>Baseball bat</u>	Dog	Tennis Racquet
	Irregular	<i>Vignette</i>	Bugs Bunny was at school. He [said / <u>sent</u>] a <i>letter</i> , which was fun.			
	Polyseme	<i>Targets</i>	Capital letter	<u>Posted letter</u>	Number	Parcel
	Regular	<i>Vignette</i>	Kermit was in the country. He [cooked / <u>heard</u>] a <i>turkey</i> , which was cool.			
	Polyseme	<i>Targets</i>	Food Turkey	<u>Animal Turkey</u>	Carrot	Bird
	Homophone	<i>Vignette</i>	Snoopy was [reading about animals / <u>watching sports</u>]. The <i>bat</i> was big.			
		<i>Targets</i>	Animal bat	<u>Baseball bat</u>	Dog	Tennis Racquet
Prior- sentence Context	Irregular	<i>Vignette</i>	Bugs Bunny was [at school / <u>the post office</u>]. The <i>letter</i> was fun.			
	Polyseme	<i>Targets</i>	Capital letter	<u>Posted letter</u>	Number	Parcel
	Regular	<i>Vignette</i>	Kermit was [having dinner / <u>in the country</u>]. The <i>turkey</i> was cool.			
	Polyseme	<i>Targets</i>	Food Turkey	<u>Animal Turkey</u>	Carrot	Bird

Table 2.

Experiment 2 Example Stimuli.

Vignette	Pictured Items (Statistical Association)			
	Plausible	Implausible	Plausible	Implausible
	Ambiguous	Ambiguous	Unambiguous	Unambiguous
Elmo watched a funny movie about a castle, and a princess, and a silly dragon. And there was a funny knight/jester.	Knight (chivalrous) (1)	Night (starry) (0)	Jester (Unambiguous)	Morning (Unambiguous)
Elmo watched a funny movie about a castle, and a princess, and a silly dragon. That was a funny night/morning.	Night (starry) (0)	Knight (chivalrous) (1)	Morning (Unambiguous)	Jester (Unambiguous)

Figure Captions

Figure 1. Two routes for resolving ambiguous words when interpreting spoken language. Use of top-down plausibility, and same-level priming between words.

Figure 2. Mean proportion of trials on which the dominant meaning was selected split by lexical ambiguity type for (a) current-sentence context and (b) prior-sentence context types. Error bars = ± 1 Standard Error of the Mean.

Figure 3. Children's accuracy identifying each sense plotted against the model's accuracy for current-sentence context (left) and prior-sentence context (right). Dashed line = line of regression.

Figure 4. (Top) Mean proportion of trials on which the dominant sense was selected by the algorithm. (Middle) Ratings given to dominant sense by adults. Numbers are mean ratings and bars are standard deviations. (Bottom) Mean proportion of trials on which the dominant sense was selected by 4-year-old children. Numbers are mean proportions selecting dominant sense and bars are 95% confidence intervals (calculated because data was binary and mean close to 0).

Appendix 1

Stimuli for Experiment 1

Target Word	Current-Distracters	Prior-Distracters (if different)	Current-sentence Context	Prior-sentence Context
Regular Polysemy				
chicken	marshmallow, goat		Barney was on vacation. He [fed/roasted] a chicken, which was fun.	Barney was [playing at a farm/at a supermarket]. The chicken was nice.
Fish	burger, crab		Oscar was at the beach. He [caught/grilled] a fish, which was exciting.	Oscar was at [the ocean/a restaurant]. The fish was exciting.
lamb	hot dog, horse		Cookie Monster was in California. He [petted/barbecued] a lamb, which was good.	Cookie Monster was [in a barn/at a barbecue]. The lamb was good.
Herb	nuts, flower		The Count was at home. He [planted/chopped] some herbs, which was messy.	The Count was in the [garden/kitchen]. The herbs were messy.
Turkey	carrot, bird		Kermit was in the country. He [heard/cooked] a turkey, which was cool.	Kermit was [in the country/having dinner]. The turkey was cool.
Tuna	ice cream, star fish		Big Bird was at the sea. He [caught/ate] some tuna, which was nice.	Big Bird was [scuba diving/at a picnic]. The tuna was nice.
Duck	pear, car	pear, turtle	Grover was in Boston. He [ran after/ate] a duck, which was good.	Grover was [swimming in a lake/at the kitchen table]. The duck was good.
shrimp	steak, shark		Miss Piggy was at the shore. She [caught/grilled] some	Miss Piggy was [snorkeling/out to lunch]. The shrimp

shrimp, which was
great. was great.

Irregular Polysemy

glasses	socks, watering can		Zoe was inside. She [put on/filled up] some glasses, which were pretty.	Zoe was getting [dressed/water]. The glasses were pretty.
letter	number, parcel		Bugs Bunny was at school. He [said/sent] a letter, which was fun.	Bugs Bunny was at [school/the post office]. The letter was fun.
button	door, zipper	head phones, zipper	Elmo was in his room. He [pushed/undid] a button, which was easy.	Elmo [turned on the music/was putting on a shirt]. The button was easy.
mouse	remote control, spider		Daffy was in his room. He [clicked/captured] a mouse, which was easy.	Daffy [needed batteries/was in the yard]. The mouse was old.
bow	hose, shoelace	campfire, box	Dora was playing. She [aimed/tied] the bow, which was difficult.	Charlie Brown was [at camp/wrapping a present]. The bow was tough
roll	cookie, handstand		Goofy was in the garden. He [chewed/performed] a roll, which was good.	Goofy was [having lunch/doing gymnastics]. The roll was good.
line	elevator, painting	elevator, lipstick	SpongeBob and Patrick were playing. They [got in/drew] a line, which was fun.	SpongeBob and Patrick were [waiting/in art class]. The line was long.
card	jacks, post-it notes		Winnie and Piglet were in the kitchen. They [played with/wrote] some cards, which was nice.	Winnie and Piglet were [playing a game/reading messages]. The cards were nice.

Homophony

bat	dog, tennis racquet		Snoopy was outside. He [chased/swung] a bat, which was big.	Snoopy was [reading about animals/watching sports]. The bat
-----	------------------------	--	---	---

				was big.
nail	picture, weed	picture, window	Dora was in her house. She [painted/pulled out] a nail, which was fun. Minnie was at the park. She [poured out/talked to] a pitcher, which was nice.	Dora was [at the beauty salon/building a house]. The nail was long. Minnie was [pouring water/playing sports]. The pitcher was nice. Elmo [heard a loud noise/was climbing]. The bark was nice.
pitcher	bucket, tennis player			
bark	bell, bench		Elmo was outside. He [heard/touched] the bark, which was ok. Dora was in her room. She [stretched/listened to] the band, which was cool. Ernie was at a party. He [lay in/chatted with] the sun, which was fun. Bert was at a dance. He [dressed up as/stayed up for] a knight, which was cool.	Dora [looked in her drawer/heard some music]. The band was cool. Ernie was [relaxing/at a family reunion]. The sun was nice Bert [was at a costume party/looked at his watch]. The knight was cool. Elmo was [in the forest/eating dessert]. The moose was cool.
band	shirt, radio	string, radio		
son/sun	bed, grandma			
knight/ night	clown, day			
moose/ mousse	bear, cake		Elmo was at camp. He [met/ate] a moose, which was cool.	

Appendix 2

Stimuli for Experiment 2

[Associated /Unassociated]

Elmo and his class were singing songs. *The teacher could play music with anything, even a band/bell. /The teacher played music with anyone, even a band/circus.*

SpongeBob, Patrick and Sandy were playing music. SpongeBob had a drum, and Patrick had a trumpet, but Sandy didn't have a guitar, *[so she had to use a band/bell. /so she had to leave the band/circus.]*

Daffy was camping in the woods. He was scared of the wild animals, *[so he yelled at the bat/blackbird. /so he brought a bat/horn.]*

Kermit was walking in a dark cave. He was nervous about the animals, [*because he saw a big bat/blackbird. /so he carried a big bat/horn.*]

Zoe was given some arrows so that she could do target practice. They were wrapped up [*along with a bow/trophy. /in a bow/cord.*]

Robin Hood aimed his arrows really well, and he won the target practice competition. He got a gold arrow tied [*to a bow/trophy. /with a bow/cord.*]

Big Bird was visiting a castle. He saw both a sword and a horse [*with the knight/jester. /during the night/morning.*]

Elmo watched a funny movie about a castle, and a princess, and a silly dragon. [*And there was a funny knight/jester. /That was a funny night/morning.*]

Dora's mom wrote a friendly note to her teacher, and then she signed [*the letter/homework. /it with a letter/number.*]

Barney was on holiday. He sent lots of postcards, and [*he only wrote a single letter/homework. /on them he wrote a single letter/number.*]

Ernie saw a little animal on his desk. It was chewing fast [*like a mouse/chipmunk. /on his mouse/apple.*]

A little animal had made a house on Piglet's desk. Pooh saw that it was [*Piglet's mouse/chipmunk. /in Piglet's mouse/apple.*]