

Disfluencies signal reference to novel objects for adults but not children*

Sarah J. OWENS, Justine M. THACKER, AND
Susan A. GRAHAM
University of Calgary, Canada

(Received 18 December 2016 – Revised 2 June 2017 – Accepted 9 August 2017)

ABSTRACT

Speech disfluencies can guide the ways in which listeners interpret spoken language. Here, we examined whether three-year-olds, five-year-olds, and adults use filled pauses to anticipate that a speaker is likely to refer to a novel object. Across three experiments, participants were presented with pairs of novel and familiar objects and heard a speaker refer to one of the objects using a fluent (“*Look at the ball/lep!*”) or disfluent (“*Look at thee uh ball/lep!*”) expression. The salience of the speaker’s unfamiliarity with the novel referents, and the way in which the speaker referred to the novel referents (i.e., a noun vs. a description) varied across experiments. Three- and five-year-olds successfully identified familiar and novel targets, but only adults’ looking patterns reflected increased looks to novel objects in the presence of a disfluency. Together, these findings demonstrate that adults, but not young children, use filled pauses to anticipate reference to novel objects.

Consider the following utterance: while at the zoo, your friend says: “*Oh! Look over there, it’s a, uh, wallaby!*” Natural speech is often disfluent, including repetitions of sounds or words, self-corrections, silent pauses, and, as in this example, filled pauses (i.e., *um, uh*). Although once

[*] This work was supported by funds from the Canada Foundation for Innovation, the Canada Research Chairs program, and the University of Calgary, and by an operating grant from the Social Sciences and Humanities Research Council of Canada (SSHRC) awarded to Susan Graham. Sarah Owens was supported by a graduate fellowship from SSHRC, and funding from Alberta Innovates Health Solutions. Justine Thacker was supported by a graduate fellowship from SSHRC. We thank Melanie Khu, Vanessa Schell, Nicole Henezi, and Charlene Parker for their assistance with this research. Some of the data presented were included in Sarah Owens’ dissertation, submitted to the University of Calgary. Address for correspondence: S. Graham, Dept. of Psychology, University of Calgary, Calgary AB, T2N 1N4, Canada; e-mail: susan.graham@ucalgary.ca

considered an unnecessary divergence from ideal speech, there is substantive evidence that speech disfluencies are produced in consistent and reliable ways within the speech stream, most often in situations which involve increased cognitive demands for a speaker. Due to regular co-occurrences in spoken language, speech disfluencies can guide the ways in which listeners process and interpret language, even helping listeners to anticipate a speaker's referential intentions (e.g., Arnold, Fagnano, & Tanenhaus, 2003; Arnold & Tanenhaus, 2011). For example, the filled pause in the example above could signal that your friend is unlikely to refer to a common or familiar animal, such as a lion or monkey. In the present experiments, we explored whether children and adults use disfluencies to infer that a speaker is about to refer to a novel rather than familiar entity.

Disfluencies are prevalent in natural language, with estimates indicating a rate of approximately six disfluencies per every 100 words spoken by adults (Fox Tree, 1995). Interjections, or filled pauses, are the most common form of disfluency to occur in the everyday speech of adults, occurring at a rate of approximately 2.56 per 100 words (Bortfeld, Leon, Bloom, Schober, & Brennan, 2001). Research has provided insight into the various contexts in which speakers have a higher probability of being disfluent. Specifically, adult speakers are more likely to be disfluent when referring to an entity or topic that is previously unmentioned within a discourse context (e.g., Arnold & Tanenhaus, 2011; Fraundorf & Watson, 2011). Disfluencies also occur more frequently in adult speech before and during clauses that are longer and more complex to generate (e.g., Arnold & Tanenhaus, 2011; Clark & Wasow, 1998; Oviatt, 1995). The number of disfluencies in children's speech also increases as the length and complexity of an utterance increases, from as early as age 2;6 (e.g., Rispoli & Hadley, 2001; Yaruss, Newman, & Flora, 1999; Zackheim & Conture, 2003). Disfluencies are also more common in situations characterized by speaker uncertainty (e.g., Krahmer & Swerts, 2005; Schober, Conrad, Dijkstra, & Ongena, 2012; Smith & Clark, 1993), although this finding is more robust for adult than child speakers (Krahmer & Swerts, 2005). Finally, for adult speakers, disfluencies occur more frequently in situations that are characterized by a greater degree of choice in referring expression (e.g., Hartsuiker & Notebaert, 2010; Schacter, Christenfeld, Ravina, & Bilous, 1991).

Due to these regular appearances in spoken language, adult listeners can use disfluencies to generate expectations about a speaker's referential intentions (e.g., Arnold *et al.*, 2003; Arnold & Tanenhaus, 2011). In particular, disfluencies lead listeners to anticipate reference to particular types of entities, including those that are new to a discourse (e.g., Arnold, Hudson Kam & Tanenhaus, 2007) and those that are novel or unfamiliar (Arnold, Tanenhaus, Altmann, & Fagnano, 2004). In one such study, Arnold and colleagues (2007) tracked adults' eye-movements as they

viewed a series of arrays of four objects on a screen, two of which were matched familiar objects that differed only in colour (e.g., a red ice cream and a black ice cream), and two of which were matched unfamiliar objects that also differed only in colour (e.g., a red squiggly shape and a black squiggly shape). Participants heard either a fluent or disfluent instruction to click on one of the objects in each array (e.g., “Click on [*the/thee uh*] red X.”). On disfluent trials, participants looked at the unfamiliar objects significantly more than they looked at the familiar objects, beginning at the time of onset of the colour word, whereas they were just as likely to look at the unfamiliar and familiar objects on fluent trials.

On the basis of findings such as those reviewed above, it has been argued that adult listeners can make situation-specific inferences about reasons for a speaker’s production difficulties, and form referential expectations accordingly. That is, in some situations, upon hearing a disfluency, adult listeners attribute the presence of the disfluency to production difficulties, make hypotheses regarding the most likely reason for the production difficulties (e.g., introducing something new to the discourse, naming an object that is unfamiliar), and consequently anticipate reference to the entity most likely to have caused the speaker to hesitate. This attributional account of disfluencies is supported by research demonstrating that the referential expectations that adult listeners form on the basis of disfluencies can be flexible (e.g., Arnold *et al.*, 2007; Barr & Seyfeddinipur, 2010; Heller, Arnold, Klein, & Tanenhaus, 2015; and see Arnold *et al.*, 2007, and Heller *et al.*, 2015, for caveats to this account). For example, Arnold and colleagues (2007) demonstrated that adult listeners could suspend their use of the disfluency cue if there was another explanation for the disruption in speech. Specifically, when listeners were told that a speaker suffered from object agnosia (a condition causing the inability to name simple objects), the listeners’ expectation for reference to unfamiliar objects was suspended and participants considered both familiar and novel objects equally for the fluent and disfluent instructions. Similarly, Barr and Seyfeddinipur (2010) demonstrated that the presence of a disfluency led adult listeners to anticipate reference to discourse-new referents, but only when those referents were new for the speaker. Taken together, this research suggests that the ways in which adult listeners respond to disfluencies often depend on their attributions about the speaker’s perspective and thought processes.

Following from research on the facilitative role of disfluencies in adults’ language processing, evidence has emerged documenting the role of disfluencies in guiding children’s referential communication as well (Kidd, White, & Aslin, 2011; Orena & White, 2015; Owens & Graham, 2016). In the first study to address this question, Kidd *et al.* (2011) demonstrated that children between 2;4 and 2;8 showed an increase in looks to objects

that were both novel and discourse-new when an instruction contained a disfluency, compared to when the instruction was spoken fluently. This effect emerged during children's second year, as children younger than 2;0 failed to reliably use disfluency as a cue to reference. Importantly, children's sensitivity to disfluencies, like that of adults, appears to reflect speaker-specific inferences about the reason for the disfluency. In a study by Orena and White (2015), three-year-olds were introduced to either a knowledgeable or a forgetful speaker, who produced fluent and disfluent utterances. Children who heard the knowledgeable speaker looked preferentially at novel, discourse-new objects during disfluent utterances, consistent with the results of the study by Kidd *et al.* (2011). In contrast, children presented with the forgetful speaker did not show the same effect – there was no difference in their looking patterns for disfluent versus fluent trials. Thus, like adults, three-year-olds can suspend their use of a disfluency when production difficulties are more likely to reflect speaker-specific characteristics than information about the objects to which a speaker is referring.

Thus, by their third year, children demonstrate an emerging ability to use filled pauses to facilitate the processing of referential communication. One question that arises, however, surrounds the types of referents that figure into children's disfluency-based predictions. Recall that disfluencies lead adults to anticipate both discourse-new and novel referents individually, consistent with a situation-dependent, flexible use of this cue (Arnold & Tanenhaus, 2011). Developmental research to date, however, has only begun to examine the types of referents children consider when hearing disfluent speech. That is, results from Kidd *et al.* (2011) and Orena and White (2015) demonstrated that filled pauses led preschoolers to anticipate reference to objects that are both discourse-new and novel, leaving it unclear which factor (discourse status or object novelty) was driving children's inferences. To better understand the communicative function of disfluencies in childhood, however, the roles of discourse status and object familiarity need to be disentangled. Certainly, appreciation of each of these factors independently would suggest a more nuanced and well-developed use of this cue. In a recent step towards this goal, Owens and Graham (2016) demonstrated that children as young as 2;6 will also anticipate reference to objects that are discourse-new on the basis of a disfluency, even when those objects are familiar to the listener.

In the present studies, we sought to examine whether disfluencies are informative for young listeners when presented with novel (i.e., unfamiliar) versus familiar referents and discourse status is held constant. Across a series of three experiments, children's (Experiments 1 and 2) and adults' (Experiment 3) processing of referential utterances was examined to

determine whether a disfluency (the filled pause [uh] in this case) would lead listeners to anticipate reference to an object which was unfamiliar. In each experiment, participants were presented with pairs of novel and familiar objects and heard a speaker refer to one of the objects using a fluent (“*Look at the ball/lep!*”) or disfluent (“*Look at thee uh ball/lep!*”) expression. The salience of the speaker’s unfamiliarity with the novel referents, and the way in which the speaker referred to the novel referents (i.e., with a noun or a description) varied across experiments.

Experiment 1

The purpose of Experiment 1 was to examine whether three- and five-year-olds use filled pauses to signal that a speaker is about to refer to a novel rather than a familiar object. We included children from two age groups to capture possible developmental progression in children’s understanding of filled pauses. There is evidence that the production of filled pauses increases as a function of age during the early years of development (Haynes & Hood, 1977; Kidd *et al.*, 2009). Thus, it is conceivable that developmental differences in attention to filled pauses may emerge as children begin to use more filled pauses in their own speech, as well as gain experience with language more generally.

Children were presented with pairs of objects, one familiar and one novel, in a referential communication task based on the procedure used by Kidd *et al.* (2011). As children viewed the object pairs on a large screen, they heard a series of three recorded utterances. See Figure 1 for an overview of the procedure. The first two utterances served to direct children’s attention to the objects, without specific reference to either object (e.g., “*I see these ones!*”). The third and final utterance was the critical utterance. Using a within-subjects design, two factors were varied across the critical utterance: (1) whether the utterance referred to the familiar or the novel object; and (2) whether the utterance was spoken fluently (“*Look at the ball/lep!*”) or disfluently (“*Look at thee uh ball/lep!*”). Our predictions focus on whether children show distinctive patterns of looking during disfluent trials compared to fluent trials. If children associate disfluencies with entities that are unfamiliar, hearing a filled pause should lead them to direct a greater proportion of looks to the unfamiliar referent, relative to when utterances are produced fluently.

Method

Participants. Data from 29 three-year-olds (17 males; $M_{age} = 3;6$, $SD = 2.02$ months) and 28 five-year-olds (15 males; $M_{age} = 5;7$, $SD = 3.13$ months) were included in the final sample. An additional 21 children were tested but

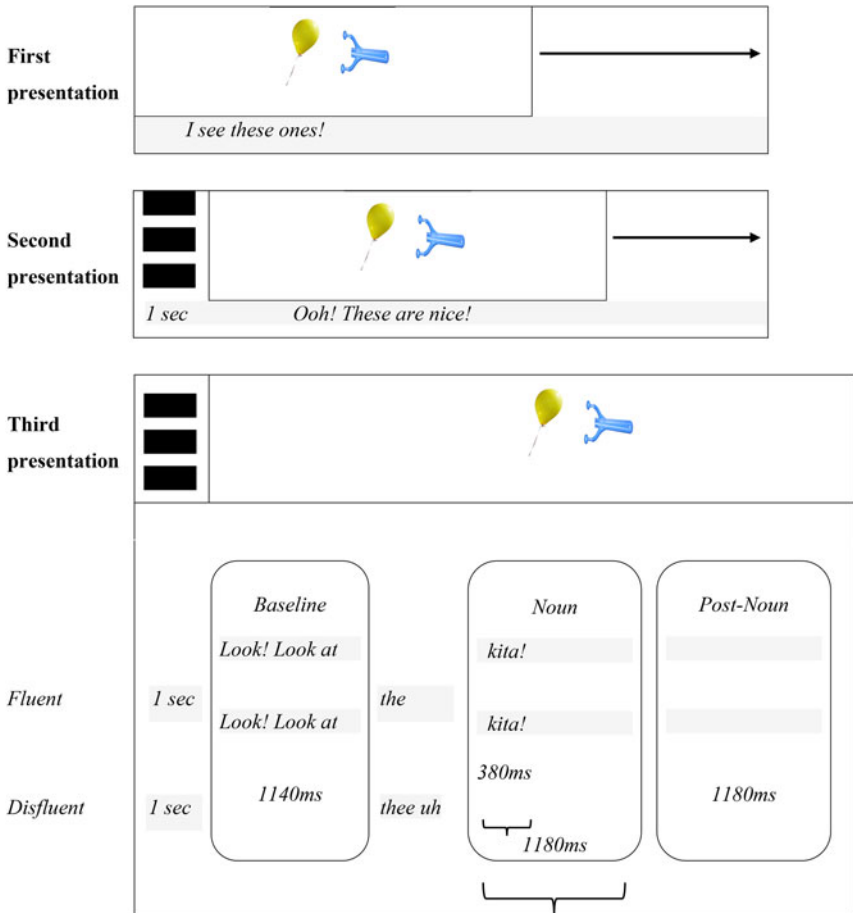


Fig. 1. Overview of a novel word test trial in Experiment 1. For test trials, each of the object pairs was presented three times in succession. The third presentation (i.e., test trial) contained the critical instruction, either fluent or disfluent, for the child to look at either the familiar or novel object.

excluded from analyses due to the following reasons: parent-reported speech delay ($n=1$), uncooperativeness or failure to complete the task ($n=6$), technical error ($n=2$), and zero looks registered to either object in the array during one or more time periods in the noun region, for all trials of a given trial type ($n=12$). Children were primarily Caucasian, from socioeconomic backgrounds that varied broadly within the more general middle class (although the latter was not formally assessed), and English was reported to be the primary spoken language.

Materials. The visual stimuli consisted of 17 photographed images of familiar objects and 17 photographed images of novel objects, which were combined to make 17 object pairs. One of the familiar–novel object pairs was used for the warm-up phase, to allow for a relevant demonstration and explanation prior to the testing phase. The remaining 16 object pairs were included in the test phase. The left–right placement of each object within a given pair was fixed, and arranged so that familiar and novel objects were presented equally as often on the left and right sides of the screen for the test trials.

Auditory stimuli consisted of utterances recorded by a female native English speaker in a soundproof booth. The speaker recorded the following two utterances, which were used to introduce all object pairs on the test trials: “*I see these ones!*” and “*Ooh, these are nice!*” The speaker also recorded two versions of the critical instruction to look at a particular object: “*Look! Look at the X!*” (fluent) and “*Look! Look at thee, uh, X!*” (disfluent), as well as an instruction to point to the target, “*Point to the X.*” The disfluent versions of the critical utterances were characterized by an unreduced pronunciation of the determiner [*the*], pronounced as [*thee*], followed by the filler [*uh*] and a brief pause. For fluent utterances, the determiner [*the*] was pronounced as [*thuh*], and the utterance contained no silent or filled pauses. Novel objects were referred to using novel nouns, which always differed in onset, and were matched in number of syllables to the name of the familiar object in the pair. To ensure consistency across critical trials, segments of the critical utterance were standardized for all test trials. Specifically, the initial prompt “*Look!*” was excised from the same recording and spliced into all fluent and disfluent critical utterances. The same recording of the disfluent determiner (“*thee, uh*”) was also excised and spliced into all disfluent critical utterances.

Apparatus. Participants’ eye fixations were tracked using a TOBII x50 remote eye-tracker connected to a 46-inch computer monitor. We recorded participants’ eye-movements to two areas of interest (AOIs) which corresponded to the location of the two objects. Gaze data were logged every 20 ms and fixations to particular AOIs were registered as looks if they were longer than 100 ms. We integrated the gaze data with the prerecorded utterances using Eye-gaze Language Integration (ELIA) software (Berman, Khu, Graham, & Graham, 2013). Calibration of the eye-tracker was conducted using Clearview software, which required children to fixate on a series of five dots, located in each corner and in the middle of the screen. A high-definition video camera was positioned behind the participants to record the experiment.

Procedure

Introduction and warm-up. Children were seated on a small chair facing the computer monitor. Once seated, the calibration procedure commenced. Accurate calibration was required on at least three out of the five test fixation points prior to initiating the experiment. Calibration for all five fixation points was achieved for 88% of participants.

Following calibration, the experimenter followed a predetermined dialogue to introduce the task. This dialogue emphasized that the speaker (“Mel”) had just found a box of toys and that while she knew the name of some of the toys, she was not familiar with all of the objects. This was followed by a warm-up trial in which children viewed a familiar–novel object pair on the screen (i.e., an airplane and a wire egg holder) and were asked by the experimenter to guess which of the two objects would be known to the speaker. The goal of the warm-up exercise was to demonstrate that the objects that would be unfamiliar to the child (i.e., the egg holder) were also unfamiliar to the speaker and experimenter. Participants received feedback on their response, such that if they identified the familiar object as the one that the speaker was likely to know, the experimenter reinforced their response (e.g., “*That’s right! Mel knows the airplane, but she’s not so sure about that one!*”). If the participant incorrectly identified the novel object as the one that Mel was likely to know, the experimenter corrected the response (e.g., “*Actually, Mel knows the airplane, but she’s not so sure about that one!*”).

Test phase. Immediately following the warm-up trial, each participant viewed a series of 16 test trials. See [Figure 1](#) for an overview of the procedure. For all trials, a pair of images was presented on a large display screen three times in succession, with a 1-second pause and blank screen between each presentation. During the first two presentations, participants heard general statements directing their attention to the screen (i.e., “*I see these ones!*” and “*Ooh, these are nice!*”). Neither object on the screen was labelled during these presentations. The third and final presentation included the critical instruction, spoken fluently or disfluently, to locate either the familiar or the novel object (“*Look! Look at [the/thee uh] X! Point to the X.*”).

The critical instruction was fluent for eight of the test trials and disfluent for the other eight test trials. For both fluent and disfluent trial types, the target was the familiar object for half of the trials, and the novel object for the other half. Thus, each participant viewed four trials of each trial type: Fluent/Familiar, Fluent/Novel, Disfluent/Familiar, and Disfluent/Novel. The 16 test trials were presented in a quasi-randomized fashion, and the target object had an equal likelihood of being located on the left and right

sides of the display. The object pairs associated with each of the four trial types were counterbalanced across participants.

Results

Children's looking patterns were analyzed during two distinct intervals (baseline and noun) of the critical utterance (“*Look! Look at [the/thee uh] X!*”), described in detail below (see Figure 1). See Figure 2 for three- and five-year-olds' proportion of looks to the novel object during each interval, as a function of fluency. The timecourse of children's proportion of looks to the novel object over the course of the critical utterance is presented in Figure 3. To analyse the proportion data, the data were first transformed using a logit transformation as this is the transformation recommended when analysing proportion data (Warton & Hui, 2011). For ease of interpretation, we report the untransformed proportions in the text and in the figures.

Baseline interval. Participants' eye-gaze during a baseline interval was analysed to examine whether looking patterns reflected any pre-existing expectations about reference, prior to the fluency manipulation. A 1140 ms time segment beginning with the onset of the critical utterance was analyzed, as this captured the portion of the utterance prior to the determiner (“*Look! Look at ...*”). The interval length, 1140 ms, represents the shortest duration of this utterance segment across trials and thus never included any fluency information.

The average proportion of looks to the novel object out of total looks to the novel and familiar objects was calculated and then submitted to the logit transformation as noted above. As listeners had not yet been exposed to the fluency manipulation or to noun information, we expected that proportion of looks to the novel object would not differ as a function of the ultimate fluency of the utterance. Results of a 2 (Age Group: Three-year-olds vs. Five-year-olds) \times 2 (Fluency: Fluent vs. Disfluent) mixed-factor ANOVA indicated no significant effects of age group ($p = .76$), fluency ($p = .42$), or an age group by fluency interaction ($p = .51$). A one-sample t -test using the untransformed proportions confirmed that the proportion of looks to the novel object ($M = 0.52$, $SD = 0.12$) did not differ from chance levels (.50) ($p = .16$). Together, these analyses indicate that children's looking patterns during baseline did not vary as a function of age group or fluency and were not indicative of an initial preference for either object.

Noun interval. If children appreciate that a filled pause can signal upcoming reference to novel objects, we expected differential looking



Experiment 1: 3-year-olds**Experiment 1: 5-year-olds**

Fig. 2. Experiment 1: three- and five-year-olds' proportion of fixations to the novel object during the baseline and noun intervals as a function of fluency. Error bars are standard error.

DISFLUENCY AS A CUE TO NOVELTY

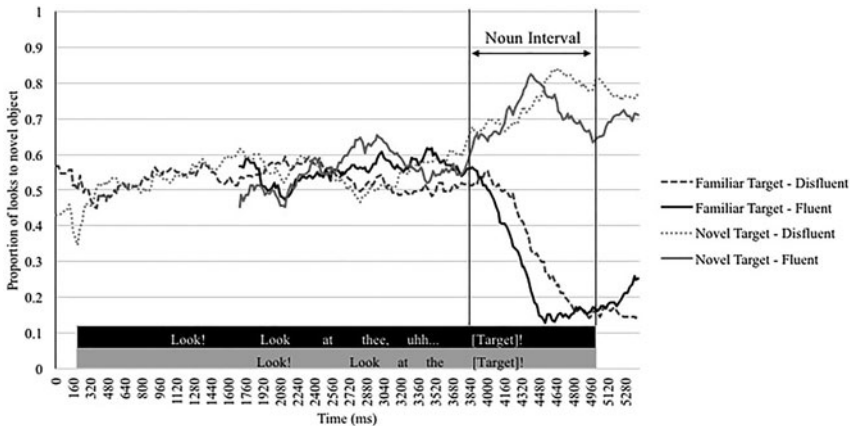


Fig. 3. Proportion of looks to the novel object over the course of the critical utterance for trials with familiar targets and trials with novel targets. During the 1180 ms window of analysis, children’s proportional looking to the novel object did not differ significantly during fluent vs. disfluent trials.

patterns for fluent and disfluent trials beginning with the determiner. As the discrepancy between the average duration of the fluent (137 ms) and disfluent (1608 ms) determiners made it unfeasible to compare children’s looks during this region, we analysed children’s looks immediately following the determiner, beginning with the onset of the critical noun. A time segment of 1120 ms, which corresponds to the length of the longest noun across trials, was analysed. We calculated the average proportion of looks to the novel object, out of the total looks to both objects, and compared these proportions for fluent and disfluent trials. As noun information was unfolding during this interval, target type (i.e., familiar or novel) was also included as a variable. Studies have demonstrated that the mutual exclusivity bias, or the tendency for children to map novel words to novel objects, emerges some time during the second year of life (e.g., Evey & Merriman, 1998; Graham, Nilsen, Collins, & Olineck, 2010; Graham, Poulin-Dubois, & Baker, 1998). Therefore, we expected that children would correctly interpret the novel label as referring to the novel object in the pair, and would arrive at the appropriate target regardless of familiar–novel status.

If the presence of a filled pause leads children to anticipate reference to a novel object, we expected a greater average proportion of looks to the novel object when the critical utterance was disfluent, compared to when it was fluent. Due to the fact that noun information was unfolding during the critical interval, we also expected children to show a greater proportion of looks to the target object. A 2 (Age Group: Three-year-olds vs.

Five-year-olds) \times 2 (Fluency: Fluent vs. Disfluent) \times 2 (Target Type: Familiar vs. Novel) mixed-factor ANOVA on the logit transformed proportions revealed a main effect of target type only ($F(1,55) = 24.21$, $\eta_p^2 = 0.31$, $p < .0001$). Children from both age groups showed a greater average proportion of looks to the novel object when it was the target ($M = 0.65$, $SD = 0.13$), compared to when the familiar object was the target ($M = 0.37$, $SD = 0.16$). One-sample t -tests indicated that the proportions were significantly above and below chance level (50%) ($t(56) = 8.54$, $p < .001$, and $t(56) = 5.98$, $p < .001$, respectively). This indicated that three- and five-year-olds successfully located the target object during the noun interval, regardless of whether the target was novel or familiar. There were no main effects of age group ($p = .81$) or fluency ($p = .94$), and no significant interaction effects (all $ps > .71$). Contrary to our hypothesis, disfluent instructions did not result in an increased proportion of looks to the novel object for three- or five-year-olds.

Although the above analyses indicated no effects of fluency, it is possible that any effects of the disfluency were masked by more salient effects of hearing the target noun label, given that we analysed looks spanning the entire noun region. That is, isolating the earliest looks within the noun interval may more clearly reflect children's use of fluency information. We therefore analysed looks during the first third of the noun interval, which was equal to 380 ms (i.e., the EARLY NOUN interval). Given that adults require approximately 200 ms to program an eye-movement in response to linguistic input (Allopenna, Magnuson, & Tanenhaus, 1998), and that a slightly longer latency may be required for children (Trueswell, 2008), we expected the eye-movement data from this interval to reflect children's processing of the utterance only up to and including the determiner. Results of a 2 (Age Group: Three-year-olds vs. Five-year-olds) \times 2 (Fluency: Fluent vs. Disfluent) \times 2 (Target Type: Familiar vs. Novel) mixed-factors ANOVA on the logit transformed proportion data yielded no significant main effects of age group, fluency, target type, nor any interactions ($ps > .24$).

Discussion

Results indicate that three- and five-year-old children failed to anticipate reference to novel objects upon hearing the speaker use disfluent referring expressions. Although it is possible that children simply do not associate filled pauses with reference to unfamiliar objects by 5;0, at least two alternative explanations are possible. First, hearing the speaker name the novel objects may have interfered with children's conceptualization of the novel objects as unfamiliar, and therefore as responsible for speaker production issues. That is, by naming the novel objects with novel nouns,

the speaker provided evidence contrary to the cover story in the introduction that the speaker was unfamiliar with those objects. A second potential explanation rests on the notion that research has yet to disentangle the effects of object novelty from those of object complexity, as they pertain to the processing of disfluent speech. That is, it is possible that the tendency for speakers to become disfluent when speaking about novel entities arises because these objects are less easily named and/or because they require lengthier descriptions. Research has demonstrated that disfluencies have a higher probability of occurring during longer utterances (Oviatt, 1995; Shriberg, 1996), and even familiar objects that require longer descriptions are associated with disfluent speech (Arnold & Tanenhaus, 2011). However, in Experiment 1, the speaker produced concise noun phrases (e.g., *the fep*), even for objects with which she was supposedly not familiar. To address these possibilities, we conducted Experiment 2.

Experiment 2

In Experiment 2, the speaker's varying degrees of familiarity with the objects in the arrays was more clearly demonstrated. Specifically, as the speaker was alleged to be unfamiliar with the novel referents, one might expect her lack of experience with those objects to result in the production of a description, rather than a definite noun phrase, when referring to the novel objects. Thus, reference to the novel objects for all critical trials was made using a description (e.g., *thing that's blue with two handles*) rather than a novel noun. See Figure 4 for an overview of the procedure.

Method

Participants. Data from 23 three-year-olds (11 males; $M_{age} = 3;6$, $SD = 1.28$ months) and 28 five-year-olds (19 males; $M_{age} = 5;3$, $SD = 2.04$ months), recruited from the same population as in Experiment 1, were included in the final sample. An additional 20 children were tested but excluded from analyses due to the following reasons: failure to complete the task ($n = 4$), technical error ($n = 6$), English was not the primary language in the child's home ($n = 2$), and zero looks were registered to either object in the array during the critical noun region, for all trials of a given trial type ($n = 8$).

Materials. The visual stimuli consisted of 20 photographs of familiar objects and 20 photographs of novel objects, combined to make 20 familiar–novel object pairs. One of the object pairs was presented to participants as part of the introduction/warm-up phase. The remaining 19 pairs were presented during the test phase: 16 of the object pairs were

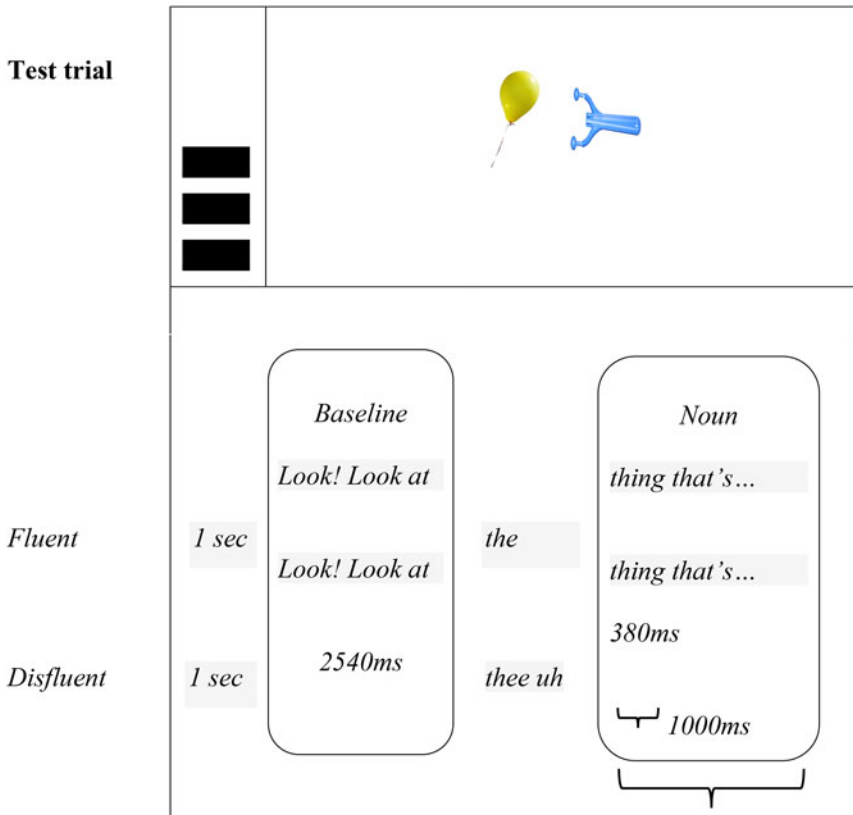


Fig. 4. Example of a novel description test trial in Experiment 2. For test trials, each of the object pairs was presented once, accompanied by the critical instruction, either fluent or disfluent, for the child to look at either the familiar or novel object. Novel objects were described rather than labelled, while familiar objects were labelled.

used for the test trials, and 3 were included as filler trials. The left–right placement of each object within a given pair was fixed, and arranged to ensure that familiar and novel objects were presented equally as often on the left and right sides of the screen for the test trials.

The auditory stimuli consisted of utterances recorded by a female native English speaker in a soundproof booth. For each of the objects included in the critical test trials, the speaker recorded a fluent and disfluent version of the critical instruction to look at the object (“*Look! Look at the X!*” and “*Look! Look at thee, uh, X!*”), as well as an instruction to point to each object (“*Point to the X!*”). The speaker referred to the familiar objects using a definite noun phrase (e.g., “*the balloon!*”), but referred to the novel

objects with a description of the image, rather than a definite noun phrase. For example, a novel object may have been described as “*the thing that’s blue with two handles*” or “*the one with the little stars*”. For the filler trials, the speaker recorded only one instruction per object pair: a disfluent instruction to look at the novel object. For the filler trial recordings, the novel objects were referred to using novel nouns. To ensure consistency across test trials, particular segments of the recorded utterances were standardized. Specifically, the initial “*Look!*” was excised from the same recording and then spliced into all test and filler trials. In addition, the same disfluent portion of the utterance, “*Look at thee, uh ...*”, was excised and spliced into all disfluent critical trials and filler trials.

Apparatus. Same as in Experiment 1.

Procedure

Introduction and warm-up. The set-up and calibration procedures were the same as in Experiment 1, with 92% of participants achieving calibration on all five points. The experimenter then followed a predetermined dialogue to introduce the task. The experimenter explained that her friend, “Mel”, had just found a box of toys, and that while Mel knew some of the toys, other toys she had never seen before. Following the introduction, children participated in a brief warm-up phase in which they viewed a familiar–novel object pair on the screen. See Figure 4 for an overview of the procedure. The experimenter clarified that one of the objects was familiar by pointing and labeling the object (“*There’s the cat ...*”). The experimenter then highlighted the novelty of the other object by describing rather than labeling the object (“*... and this white round thing*”), and explicitly stated that neither she nor her friend Mel knew what to call the novel object (“*Mel doesn’t know what to call that one, and neither do I*”). Participants were provided with both verbal and non-verbal cues to emphasize the unfamiliarity of the novel object in the warm-up pair (i.e., the experimenter stated that she did not know the name of the novel object while shrugging her shoulders and shaking her head from side to side). All participants were included, regardless of their response to the familiarization question.

Test phase. There were two primary differences between the test phase procedure for Experiment 2 and the test phase procedures for Experiment 1. First, each object pair was presented a single time instead of three times, to reduce the risk of participant fatigue. Second, filler trials were included to prevent participants from developing expectations about the way in which novel objects would be labelled (i.e., to prevent the

expectation that the speaker would always describe, rather than label, the novel objects).

For all trials in the test phase, a pair of images was presented a single time on a large display screen. With each object pair presentation, children heard a prerecorded utterance pertaining to one of the objects on the screen. Specifically, participants heard a critical instruction to look at one of the objects, followed by an instruction to point to the target (“*Look! Look at [the/thee, uh] X! Point to the X.*”). As in Experiment 1, critical utterances were either fluent or disfluent, and made reference to either the familiar or the novel object. Thus, participants viewed four critical test trials of each trial type, Fluent/Familiar, Fluent/Novel, Disfluent/Familiar, and Disfluent/Novel, presented in a quasi-randomized fashion. Recall that, for the test trials in Experiment 2, all of the novel objects were referred to using descriptions rather than nouns.

Results

Children’s looking patterns were analysed during baseline and noun intervals as described below and depicted in Figure 4. See Figure 5 for three- and five-year-olds’ proportion of looks to the novel object during each interval, as a function of fluency. The timecourse of children’s proportion of looks to the novel object over the course of the critical utterance is presented in Figure 6.

Baseline interval. The baseline interval represented the same portion of the utterance (“*Look! Look at ...*”) as in Experiment 1, but with an overall duration of 2540 ms. This length was chosen because it represented the shortest duration of “*Look! Look at ...*” across trials for Experiment 2. Results of a 2 (Age Group: Three-year-olds vs. Five-year-olds) \times 2 (Fluency: Fluent vs. Disfluent) mixed-factor ANOVA on the logit transformed data indicated a significant main effect of age group ($p = .021$), fluency ($p = .001$), and an age group by fluency interaction ($F(1,49) = 5.67$, $\eta_p^2 = 0.10$, $p = .022$). Follow-up analyses indicated that three-year-olds looked more to the novel object during fluent utterances ($M = 0.61$, $SD = 0.09$) than during disfluent utterances ($M = 0.49$, $SD = 0.12$) ($t(22) = 3.92$, $p = .001$). In contrast, five-year-olds’ looking to the novel object during baseline did not vary as a function of the ultimate fluency of the utterance ($p = .41$). Comparisons to chance using the untransformed data indicated that three-year-olds’ looks to the novel object were above chance in the disfluent utterances ($p < .001$), but at chance in the fluent utterances ($p = .78$). Five-year-olds looked to the novel object at levels significantly greater than chance in the fluent and disfluent utterances ($ps > .01$). The early preference for the novel object is not surprising, given that children

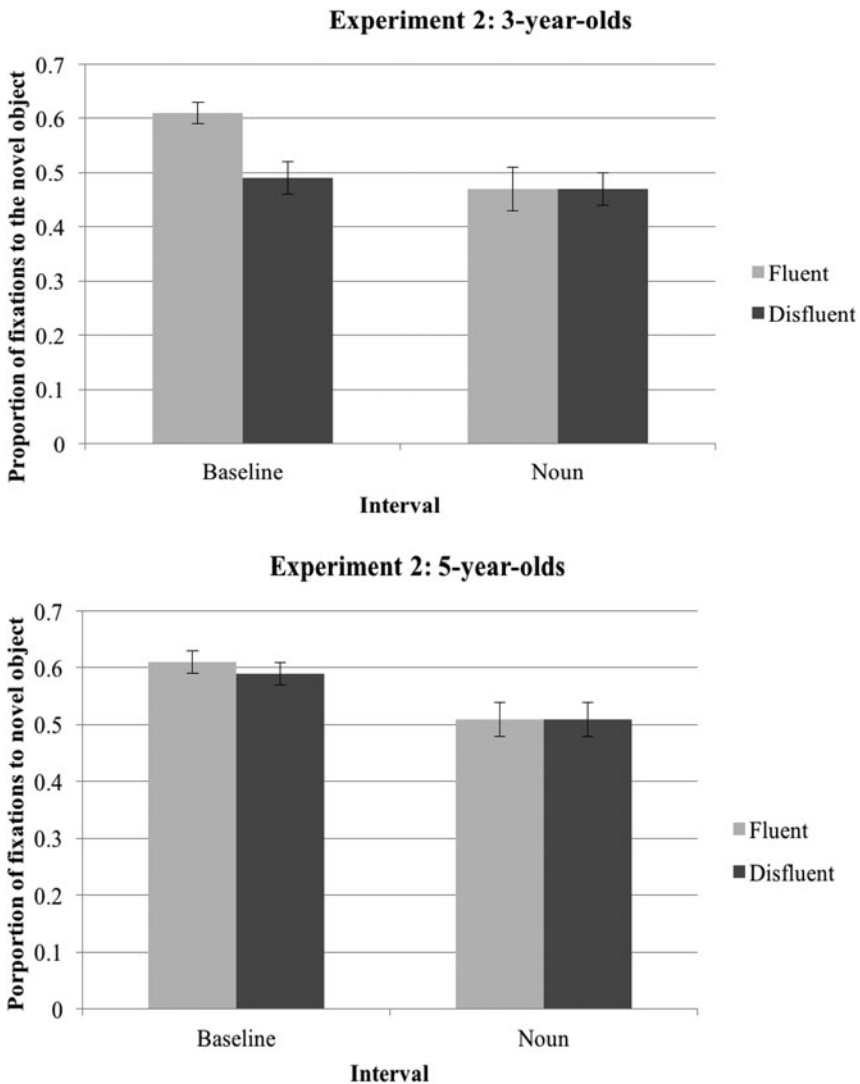


Fig. 5. Experiment 2: three- and five-year-olds' average proportion of fixations to the novel object during the baseline and noun intervals as a function of fluency. Error bars are standard error.

only viewed each array once in Experiment 2, unlike in Experiment 1. That is, the level of visual interest was likely higher for the novel objects, simply because they were objects that children had never seen before. What is more surprising is three-year-olds' lack of preference in the disfluent utterances,



Fig. 6. Proportion of looks to the novel object over the course of the critical utterance for trials with familiar targets and trials with novel targets. During the 1000 ms window of analysis, children's proportional looking to the novel object did not differ significantly during fluent vs. disfluent trials.

given that no fluency information is available during this baseline period. Critically, however, our core questions rest on comparisons between the disfluent and fluent utterances during the noun interval.

Noun interval. The noun interval began with the onset of the critical noun and corresponded to the length of the longest familiar noun across trials (1000 ms). Note that for the novel target trials, the description always began with either *one* or *thing*; these words were considered to be the critical nouns for the novel target trials. Unlike Experiment 1, the noun interval for Experiment 2 included linguistic information other than the noun for the novel target trials. The average length of the nouns *one* or *thing* was 321 ms, while the average total length of the descriptions was 2557 ms. Thus, while the noun interval for the familiar target trials only ever included looks during the critical noun and post-noun silence, the noun interval for the novel target trials included the non-specific noun *one* or *thing*, plus a portion of the description up to 1000 ms.

A 2 (Age Group: Three-year-olds vs. Five-year-olds) \times 2 (Fluency: Fluent vs. Disfluent) \times 2 (Target Type: Familiar vs. Novel) mixed-factors ANOVA revealed a main effect of target type only ($F(1,49) = 13.57$, $\eta_p^2 = 0.22$, $p < .001$). Children directed a greater average proportion of looks to the novel object when it was the target ($M = 0.58$, $SD = 0.16$) versus when the familiar object was the target ($M = 0.40$, $SD = 0.17$). One-sample *t*-tests indicated that both proportions were significantly different from chance

level (50%) ($t(50) = 3.36, p < .001$, and $t(50) = 3.90, p < .001$). There were no main effects of fluency or age group, and no interaction effects (all $ps > .13$). Thus, children successfully identified the target referent during the noun interval, with no evidence for an influence of the fluency manipulation on fixation patterns during this interval.

As in Experiment 1, we also analysed looks during the early noun interval (the first 380 ms of the unfolding description), to address the possibility of early fluency effects. Results of a 2 (Age Group: Five-year-olds vs. Three-year-olds) \times 2 (Fluency: Fluent vs. Disfluent) \times 2 (Target Type: Familiar vs. Novel) mixed-factors ANOVA yielded no significant main effects or interactions ($ps > .29$).

Discussion

The results of Experiment 2 indicated that, even when provided with an explicit description of the speaker's limited knowledge about the novel objects, as well as evidence that the speaker was unable to consistently name those objects, both three- and five-year-old children failed to show an increased proportion of looks to the novel object upon hearing a disfluency. This would seem to suggest that the presence of a filled pause does not lead children to expect reference to novel objects. An alternative explanation, however, relates to the types of novel objects used in these experiments. In research with adults, Arnold and colleagues (2007) used referents that were novel and had a complex, unconventional appearance (obscure symbols or shapes). Conversely, the novel referents presented in our experiments were not necessarily any more visually complex than the familiar objects with which they were paired, and were conventional in the sense that they would likely be perceived as nameable objects. Thus, it is possible that results from studies with adult listeners require a more nuanced interpretation, one that distinguishes between referents that are novel from referents that are less likely to have a conventional name. Perhaps adult listeners use the presence of filled pauses to predict reference to complex or unconventional, rather than novel, objects. To assess this possibility, we repeated Experiments 1 and 2 with adult participants.

Experiment 3

In Experiment 3, university students were presented with the same stimuli and procedures used in Experiments 1 and 2. If adult participants fail to show a pattern of expectation for reference to novel objects on the basis of hearing a disfluent utterance, as they have been shown to do in the existing adult literature, then the data would point to a 'nameability',

rather than novelty, explanation for the adult studies. On the other hand, if adult participants do show an increased proportion of looks to novel (but nameable) objects upon hearing a disfluency, then the data would be suggestive of important differences in the kinds of situations that young children and adults associate with filled pauses.

Method

Participants. Data from 64 university students ($M_{age} = 22;1$, $SD = 3;9$) were included in the final sample. An additional three adult participants were tested but excluded from analyses due to the following reasons: experimenter error ($n = 1$), and zero looks registered to either object in the array during the noun region, for all trials of a given trial type ($n = 2$).

Materials, apparatus, and procedure. Adults were tested in either the procedures used in Experiment 1 ($n = 34$), or Experiment 2 ($n = 30$). The majority (95%) of adults achieved calibration for all five calibration points. For ease of reference, we will refer to the adults tested in the Experiment 1 procedure as the ‘Novel Noun’ group and those tested in the Experiment 2 procedure as the ‘Novel Description’ group. The stimuli and procedures were identical to those used in the respective child experiments.

Results

Adults’ gaze data was analysed using the same intervals, with the same time-points, described in Experiments 1 and 2 (see [Figures 1](#) and [4](#)). See [Figure 7](#) for adults’ proportion of fixations to the novel object during baseline and noun intervals, as a function of fluency, and [Figure 8](#) for the timecourse of adults’ proportion of looks to the novel object, as a function of procedure, over the course of the critical utterance.

Baseline interval. Note that $n = 63$ for the baseline analyses, as one participant was removed due to zero looks registered to either object during this interval. Results of a 2 (Procedure: Novel Noun vs. Novel Description) \times 2 (Fluency: Fluent vs. Disfluent) mixed-factor ANOVA on the logit transformed data indicated no significant main effects or interactions ($p_s > .10$). One-sample t -tests indicate that adults in both groups looked to the novel object at levels significantly greater than chance ($p_s < .035$). Thus, adults in both groups show a preference for the novel object during baseline, similar to children tested in Experiment 2.

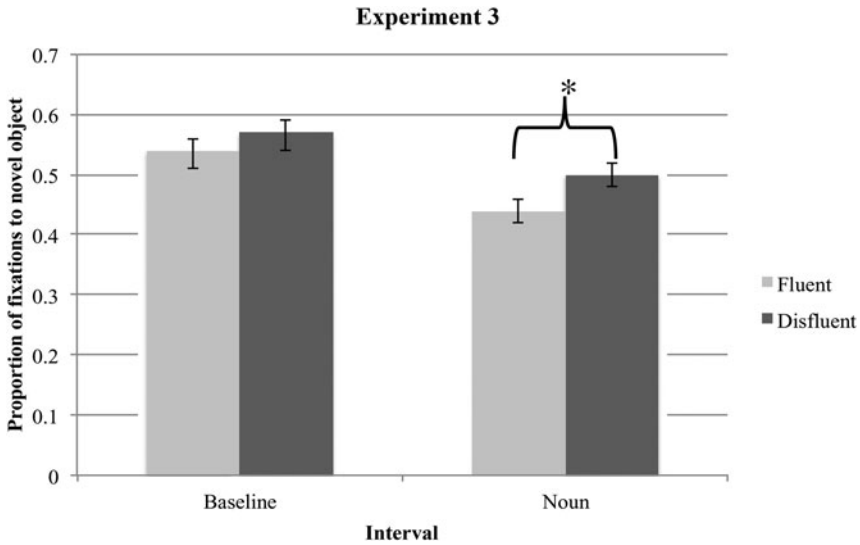


Fig. 7. Experiment 3: adults' average proportion of fixations to the novel object during the baseline and noun interval as a function of fluency. During the noun interval, adults' fixations to the novel object were significantly greater during disfluent compared to fluent utterances ($p < .05$). Error bars are standard error.

Noun interval. Results of a 2 (Procedure: Novel Noun vs. Novel Description) \times 2 (Fluency: Fluent vs. Disfluent) \times 2 (Target Type: Familiar vs. Novel) mixed-factor ANOVA on the logit transformed proportions yielded significant main effects of procedure ($F(1,62) = 5.59$, $\eta_p^2 = 0.08$, $p = .021$), target type ($F(1,62) = 39.51$, $\eta_p^2 = 0.39$, $p < .0001$), and fluency ($F(1,62) = 13.61$, $\eta_p^2 = 0.18$, $p < .0001$). No other effects reached significance ($ps > .21$).

The main effect of procedure indicated that adults looked more at the novel object in the Novel Noun procedure ($M = 0.49$, $SD = 0.11$) versus in the Novel Description procedure ($M = 0.45$, $SD = 0.08$), collapsed across fluency and target type. Critically, however, no interactions with procedure were significant ($ps > .21$). The main effect of target type indicated that adults correctly located the target during the noun interval across both procedures, as they showed a greater average proportion of looks to the novel object when it was the target ($M = 0.64$, $SD = 0.14$) versus when the familiar object was the target ($M = 0.30$, $SD = 0.12$).

Of greatest interest for our primary research question was the main effect of fluency. Unlike the children tested in Experiments 1 and 2, adults showed a greater average proportion of looks to the novel object when the utterance was disfluent ($M = 0.51$, $SD = 0.15$) versus when the utterance was fluent

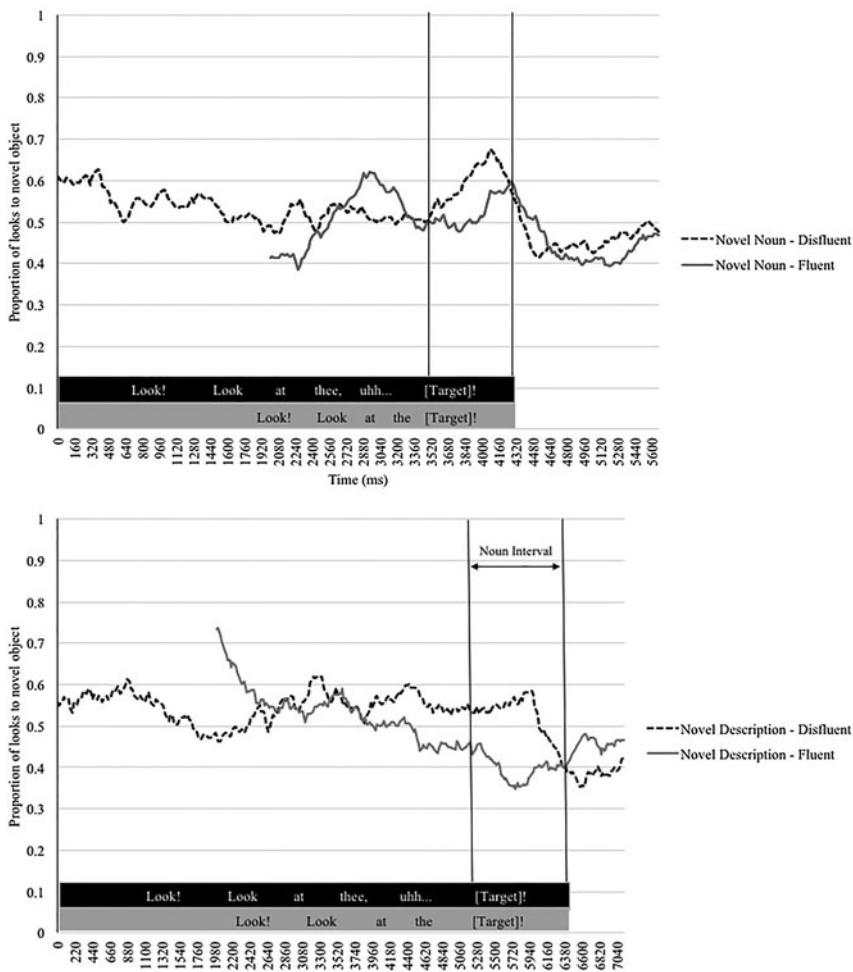


Fig. 8. Proportion of looks to the novel object over the course of the critical utterance for the Novel Noun procedure (top panel) and the Novel Description procedure (bottom panel). During the noun interval, adults' proportional looking to the novel object was significantly greater during disfluent trials compared to fluent trials ($p < .0001$).

($M = 0.44$, $SD = 0.12$). One-sample t -tests revealed that adults' average proportions of looks to the novel object were at chance (50%) for disfluent trials ($p = .51$), but significantly below chance for the fluent trials ($p < .001$).

An examination of adults' looking during the early portion of the noun, as computed in Experiments 1 and 2, further revealed only an effect of fluency ($F(1,62) = 10.94$, $\eta_p^2 = 0.15$, $p = .002$). There were no main effects of

target type or procedure, nor any interactions ($ps > .14$). Thus, prior to integrating noun or description information, adults spent a greater proportion of time looking at the novel object for disfluent trials ($M = 0.55$, $SD = 0.22$) than for fluent trials ($M = 0.44$, $SD = 0.12$). One-sample t -tests revealed that adults' average proportions of looks to the novel object were significantly above chance for disfluent trials ($p = .049$), but at chance for the fluent trials ($p = .37$).

Cross-Experiment Comparisons

In the next set of analyses, we directly compared children's and adults' performance using a 3 (Age group: Three-year-olds, Five-year-olds, Adults) $\times 2$ (Study Procedure: Novel Noun vs. Novel Description) $\times 2$ (Fluency: Fluent vs. Disfluent) $\times 2$ (Target Type: Familiar vs. Novel) mixed-factor ANOVA on the logit transformed proportions. This analysis indicated a significant main effect of target type ($F(1,166) = 70.531$, $\eta_p^2 = 0.30$, $p < .0001$). As in the analyses reported in each experiment, this main effect indicates that all participants of all ages showed a greater average proportion of looks to the novel object when it was the target versus when the familiar object was the target. Of greater interest is the significant age group by fluency interaction ($F(1,166) = 5.632$, $\eta_p^2 = 0.064$, $p = .004$). Again, mirroring the analyses reported for each individual experiment, only the adults showed a greater average proportion of looks to the novel object when the utterance was disfluent ($M = 0.51$, $SD = 0.15$) versus when the utterance was fluent ($M = 0.44$, $SD = 0.12$). No other main effects or interactions were significant ($ps > .19$). These analyses directly demonstrate that adult listeners, but not child listeners, were more likely to look at novel objects when the speaker was disfluent versus when she was fluent.

The aforementioned analyses indicate that, although adults did use filled pauses to form expectations about reference to novel objects, children showed no such differences as a function of utterance fluency. These frequentist analyses, however, do not allow us to conclude that children do not expect filled pauses to signal reference to a novel object – only that we failed to find support for the hypothesis that children would look longer at the novel object during disfluent utterances than during fluent utterances (Dienes, 2014; Dienes & Mclatchie, 2017). To more directly assess whether children DO NOT associate filled pauses with reference to novel objects, we used a Bayes factor analysis to assess the relative strength of the evidence for the alternative hypothesis (i.e., increased looking to the novel object when utterances were disfluent versus when they were fluent) versus the null hypothesis (Dienes, 2014). Using the on-line calculator associated with Dienes (2008; <www.lifesci.sussex.ac.uk/home/Zoltan_Dienes/inference/Bayes.htm>), we computed a Bayes factor for the

disfluent–fluent mean difference in looking to the novel object during the noun region (logit transformed; Mean difference = -0.28214 , SE = 0.2592). We collapsed across study protocols and modelled the alternative hypothesis (H_1) by using the relevant mean difference (using the logit transformed means) from the adult group (1.00320) as the standard deviation of a half-normal distribution. We used a half-normal or one-tailed parameter, as our prediction was directional; that is, we predicted that child participants would look more at the novel object when hearing disfluent utterances versus fluent utterances. In interpreting the resulting Bayes Factor, we followed the conventions for interpretation described by Dienes (2014), wherein a BF of 3 or above indicates substantial evidence for the alternative rather than the null hypothesis, while a BF of $1/3$ or below is taken to indicate substantial evidence for the null rather than alternative hypothesis. Our resulting Bayes Factor ($B_{H(0, 1.003)} = 0.13$) provides substantial evidence for the null hypothesis. This analysis, thus, strongly suggests that children do not associate filled pauses with reference to novel objects.

General Discussion

The starting point for these experiments was to examine whether the presence of a filled pause would signal an increased likelihood of reference to novel entities for child listeners. The results indicated that three- and five-year-old children did not direct a greater proportion of looks toward novel objects upon hearing a filled pause. Instead, children’s eye-gaze reflected their ability to quickly and correctly locate both familiar and novel target referents, with no observed effect of utterance fluency. In contrast, adults did use filled pauses to form predictions about reference to novel objects. Together, the results provide insight into the types of referents children and adults consider in their disfluency-based predictions, and point to important developmental differences in listeners’ attention to disfluencies in referential communication. We discuss these findings in turn.

Across two studies, our findings indicate that three- and five-year-olds do not anticipate reference to a novel object when hearing a disfluent utterance, even when the speaker’s limited knowledge about the novel objects is emphasized and when the speaker is unable to consistently name those objects. Instead, children’s looking patterns reflected sensitivity to linguistic input alone. That is, when hearing an instruction to locate the referent of a familiar word, children quickly identified the correct referent as the noun was unfolding. When instructed to find the referent of a novel noun (Experiment 1) or description of a novel object (Experiment 2), children also quickly landed upon the intended referent regardless of

utterance fluency, suggesting operation of a mutual exclusivity bias (Markman, 1991; Merriman, Bowman, & MacWhinney, 1989).

In contrast to the results with children, adults showed an increased proportion of looks to novel objects upon hearing a filled pause, adding to the literature on adults' sensitivity to disfluencies in two key ways. First, in previous research, Arnold *et al.* (2007) demonstrated that adults anticipated reference to objects that were novel and had a complex, unconventional appearance (obscure symbols or shapes) upon hearing a disfluent utterance. Conversely, the novel referents presented in these experiments were not necessarily any more visually complex than the familiar objects with which they were paired, and were conventional in a sense that they would likely be perceived as nameable objects. Thus, our findings demonstrate that object novelty alone is sufficient to provoke this disfluency-based anticipation in adults. Second, there was no effect of procedure in Experiment 3, meaning that the presence of a filled pause biased adults toward novel objects both when the speaker used a novel noun and when she used a longer description. This finding contributes to a recent line of research that has sought to dissociate disfluency attributions associated with object novelty versus description length. In one such study, Heller and colleagues (2015) taught adult participants a mini-artificial lexicon for a set of novel shapes, and found that disfluent instructions biased listeners toward the novel shapes with novel names, and against shapes that required a novel conceptualization and planning of a longer description. This study was the first to demonstrate that disfluency can bias listeners AGAINST objects that lack a conventional name, perhaps because the novel labels were perceived as more difficult to produce than longer descriptions for the unnamed objects. However, in another experiment, when listeners were given the opportunity to practise producing the novel names, this disfluency bias was eliminated (Heller, Gorman, & Tanenhaus, 2012). The experiments by Heller and colleagues, taken together with findings from the present study, demonstrate that adult listeners rely on situation-specific reasoning to determine which of two objects is the 'harder to name object', and thus more likely to follow production of a disfluency.

What might account for developmental differences in attention to novel objects upon hearing a filled pause? First, it is possible that children appreciate that filled pauses signal production problems in the same way that adults are thought to, but have a comparatively limited understanding of the types of situations that are likely to cause a speaker difficulty. That is, perhaps object novelty alone is simply not one of the factors that children associate with speaker production difficulties. Children have a great deal of experience with situations in which adults know, and have no difficulty producing, the names of referents that they themselves have

never encountered previously. Furthermore, research suggests that children develop awareness of their own lexical ignorance sometime between 2;6 and 4;0 (Marazita & Merriman, 2004), and by 4;0 they expect adults to know information that they themselves do not know (Taylor, Cartwright, & Bowden, 1991). It is possible then, that for children, one's own lack of familiarity with a referent does not serve as a reliable predictor of production difficulty for adult speakers. Additionally, children in these experiments had evidence that the speaker was knowledgeable about object labels, based on her accurate naming of familiar referents. Children have been shown to expect that individuals who provide accurate labels for familiar objects will have greater word knowledge in general (Brosseau-Liard & Birch, 2010), and the children in these experiments may therefore have had reason to expect the speaker to also know the names of the novel objects. One way to test this hypothesis would be to examine whether children are more likely to show evidence for the disfluency effect in the case of a child speaker naming the novel and familiar objects. Children show an appreciation that they and their same-age peers lack knowledge that adults have by the time they are 4;0 (Taylor *et al.*, 1991). Thus, it is possible that children would more readily identify object novelty as a source of production difficulties for a speaker who they perceive as being the same age as them.

Recall that Kidd *et al.* (2011) and Orena and White (2015) demonstrated that filled pauses led preschoolers to anticipate reference to objects that were both discourse-new and novel, leaving it unclear which factor (discourse status or object novelty) was driving children's inferences. To better understand the communicative function of disfluencies in childhood, the roles of discourse status and object familiarity need to be disentangled. To this end, Owens and Graham (2016) demonstrated that children as young as 2;0 will anticipate reference to familiar objects that are discourse-new, on the basis of a disfluency. The present study, however, suggests that sensitivity to the relation between disfluency and novelty develops much later. As discussed above, one possible reason for this later development is that young children may not learn about the disfluency–novelty relation through attending to adult-produced speech. If this is the case, then a subsequent question arises: If children do not learn about this relation from adults, then from whom do they learn? Three possible explanations exist. First, it may be the case that older children DO learn about the disfluency–novelty relation from adults, but that this process does not begin until they have acquired a sufficient vocabulary and are able to perceive the disfluency–novelty relation in adult-produced speech. Second, it is possible that children learn about this relation from their peers. Indeed, researchers have found that the development of children's expressive language abilities can be positively stimulated by exposure to

the language of their peers (Mashburn, Justice, Downer, & Pianta, 2009). A third possibility is that children may accrue evidence for this relation through attending to their own language production over time. It is well established that filled pauses are often produced in an effort to gain processing time (Goldman-Eisler, 1968; Rochester, 1973; Schacter *et al.*, 1991). Thus, if children are spontaneously disfluent in situations where they are unsure of or do not know which label to use for a novel object, then it follows that, over time, evidence will accumulate for the relation between filled pauses and unfamiliar objects. Future research should aim to further explore the developmental trajectory of the disfluency–novelty relation in order to elucidate which discourse-related processes are at play.

In summary, the results of the present experiments contribute to a growing body of research demonstrating that speech disfluencies, once considered irrelevant or even detrimental to speech processing, provide listeners with important information in a referential communication context. Furthermore, the findings offered here clarify previous research by showing that, while adults use filled pauses to form predictions in the context of object novelty, three- and five-year-old children do not. The development of children's understanding of disfluencies and their relation with speech production difficulties is likely to continue throughout childhood, until young listeners attain the advanced flexible approach observed in adults.

REFERENCES

- Allopenna, P. D., Magnuson, J. S., & Tanenhaus, M. K. (1998). Tracking the time course of spoken word recognition using eye-movements: evidence for continuous mapping models. *Journal of Memory and Language*, 38(4), 419–39.
- Arnold, J. E., Fagnano, M., & Tanenhaus, M. K. (2003). Disfluencies signal thee, um, new information. *Journal of Psycholinguistic Research*, 32(1), 25–36.
- Arnold, J. E., Hudson Kam, C., & Tanenhaus, M. K. (2007). If you say thee uh- you're describing something hard: the on-line attribution of disfluency during reference comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(5), 914–30.
- Arnold, J. E., & Tanenhaus, M. K. (2011). Disfluency effects in comprehension: how new information can become accessible. In E. Gibson & N. J. Pearlmuter (Eds), *The processing and acquisition of reference*, 197–217. Cambridge, MA: MIT Press.
- Arnold, J. E., Tanenhaus, M. K., Altmann, R. J., & Fagnano, M. (2004). The old and thee, uh, new: disfluency and reference resolution. *Psychological Science*, 15(9), 578–82.
- Barr, D. J., & Seyfeddinipur, M. (2010). The role of fillers in listener attributions for speaker disfluency. *Language and Cognitive Processes*, 25(4), 441–55.
- Berman, J. M. J., Khu, M., Graham, I., & Graham, S. A. (2013). ELIA: a software application for integrating spoken language and eye movements. *Behavior Research Methods*, 45(3), 646–55.
- Bortfeld, H., Leon, S. D., Bloom, J. E., Schober, M. F., & Brennan, S. E. (2001). Disfluency rates in conversation: effects of age, relationship, topic, role, and gender. *Language and Speech*, 44(2), 123–47.

- Brosseau-Liard, P. E., & Birch, S. A. (2010). 'I bet you know more and are nicer too!' What children infer from others' accuracy. *Developmental Science*, 13(5), 772–8.
- Clark, H. H., & Wasow, T. (1998). Repeating words in spontaneous speech. *Cognitive Psychology*, 37(3), 201–42.
- Dienes, Z. (2008). *Understanding psychology as a science: an introduction to scientific and statistical inference*. Basingstoke: Palgrave Macmillan.
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5, 781. doi: 10.3389/fpsyg.2014.00781
- Dienes, Z., & Mclatchie, N. (2017). Four reasons to prefer Bayesian analyses over significance testing. *Psychonomic Bulletin & Review*, 1–12. doi: 10.3758/s13423-017-1266-z
- Evey, J. A., & Merriman, W. E. (1998). The prevalence and the weakness of an early name mapping preference. *Journal of Child Language*, 25(1), 121–47.
- Fox Tree, J. E. (1995). The effects of false starts and repetitions on the processing of subsequent words in spontaneous speech. *Journal of Memory and Language*, 34(6), 709–38.
- Fraundorf, S. H., & Watson, D. G. (2011). The disfluent discourse: effects of filled pauses on recall. *Journal of Memory and Language*, 65(2), 161–75.
- Goldman-Eisler, F. (1968). *Psycholinguistics: experiments in spontaneous speech*. New York: Academic Press.
- Graham, S. A., Nilsen, E. S., & Collins, S., & Olineck., K. M. (2010). The role of gaze direction and mutual exclusivity in guiding 24-month-olds' word mappings. *British Journal of Developmental Psychology*, 28, 449–65.
- Graham, S. A., Poulin-Dubois, D., & Baker, R. K. (1998). Infants' disambiguation of novel object words. *First Language*, 18(53), 149–64.
- Hartsuiker, R. J., & Notebaert, L. (2010). Lexical access problems lead to disfluencies in speech. *Experimental Psychology*, 57, 169–77.
- Haynes, W. O., & Hood, S. B. (1977). Language and disfluency variables in normal speaking children from discrete chronological age groups. *Journal of Fluency Disorders*, 2(1), 57–74.
- Heller, D., Arnold, J. E., Klein, N., & Tanenhaus, M. K. (2015). Inferring difficulty: flexibility in the real-time processing of disfluency. *Language and Speech*, 58(2), 190–203.
- Heller, D., Gorman, K. S., & Tanenhaus, M. K. (2012). To name or to describe: shared knowledge affects referential form. *Topics in Cognitive Science*, 4(2), 290–305.
- Kidd, C., White, K. S., & Aslin, R. N. (2009). Children's use of disfluencies for pragmatic inference in lexical development. In *Proceedings of the 31st Annual Meeting of the Cognitive Science Society*, 31, 1494–9.
- Kidd, C., White, K. S., & Aslin, R. N. (2011). Toddlers use speech disfluencies to predict speakers' referential intentions. *Developmental Science*, 14(4), 925–34.
- Krahmer, E., & Swerts, M. (2005). How children and adults produce and perceive uncertainty in audiovisual speech. *Language and Speech*, 48(1), 29–53.
- Marazita, J. M., & Merriman, W. E. (2004). Young children's judgment of whether they know names for objects: the metalinguistic ability it reflects and the processes it involves. *Journal of Memory and Language*, 51(3), 458–72.
- Markman, E. M. (1991). *Categorization and naming in children: problems of induction*. Cambridge, MA: MIT Press.
- Mashburn, A. J., Justice, L. M., Downer, J. T., & Pianta, R. C. (2009). Peer effects on children's language achievement during pre-kindergarten. *Child Development*, 80(3), 686–702.
- Merriman, W. E., Bowman, L. L., & MacWhinney, B. (1989). The mutual exclusivity bias in children's word learning. *Monographs of the Society for Research in Child Development*, 54(3/4), i–129.
- Orena, A. J., & White, K. S. (2015). I forget what that's called! Children's online processing of disfluencies depends on speaker knowledge. *Child Development*, 86(6), 1701–9.
- Oviatt, S. (1995). Predicting spoken disfluencies during human-computer interaction. *Computer Speech & Language*, 9(1), 19–35.
- Owens, S., & Graham, S. A. (2016). Thee, uhh disfluency effect in preschoolers: a cue to discourse status. *British Journal of Developmental Psychology*, 34, 388–401.

- Rispoli, M., & Hadley, P. (2001). The leading-edge: the significance of sentence disruptions in the development of grammar. *Journal of Speech, Language, and Hearing Research*, 44(5), 1131–43.
- Rochester, S. R. (1973). The significance of pauses in spontaneous speech. *Journal of Psycholinguistic Research*, 2(1), 51–81.
- Schachter, S., Christenfeld, N., Ravina, B., & Bilous, F. (1991). Speech disfluency and the structure of knowledge. *Journal of Personality and Social Psychology*, 60(3), 362–7.
- Schober, M. F., Conrad, F. G., Dijkstra, W., & Ongena, Y. P. (2012). Disfluencies and gaze aversion in unreliable responses to survey questions. *Journal of Official Statistics*, 28(4), 555–82.
- Shriberg, E. (1996). Disfluencies in switchboard. In *Proceedings of International Conference on Spoken Language Processing*, 96, 11–14.
- Smith, V. L., & Clark, H. H. (1993). On the course of answering questions. *Journal of Memory and Language*, 32(1), 25–38.
- Taylor, M., Cartwright, B. S., & Bowden, T. (1991). Perspective taking and theory of mind: Do children predict interpretive diversity as a function of differences in observers' knowledge? *Child Development*, 62(6), 1334–51.
- Trueswell, J. C. (2008). Using eye movements as a developmental measure within psycholinguistics. In I. A. Sekerina, E. M. Fernandez & H. Clahsen (Eds), *Developmental psycholinguistics: on-line methods in children's language processing*, 73–96. Philadelphia, PA: John Benjamins Publishing Company.
- Warton, D. I., & Hui, F. K. C. (2011). The arcsine is asinine: the analysis of proportions in ecology. *Ecology*, 92(1), 3–10.
- Yaruss, J. S., Newman, R. M., & Flora, T. (1999). Language and disfluency in nonstuttering children's conversational speech. *Journal of Fluency Disorders*, 24(3), 185–207.
- Zackheim, C. T., & Conture, E. G. (2003). Childhood stuttering and speech disfluencies in relation to children's mean length of utterance: a preliminary study. *Journal of Fluency Disorders*, 28(2), 115–42.