Children's Selective Trust in Informants: The Role of Speech Characteristics

by

Shaneene Heupel

A thesis

presented to the University of Waterloo in fulfilment of the

thesis requirement for the degree of

Master of Arts

in

Psychology

Waterloo, Ontario, Canada, 2021

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Background. The main way children learn new information is through other people. However, not all people have information that is valuable and useful to them. Research suggests that children pay attention to many cues to help them gauge an informant's trustworthiness as an information source. These cues include the speaker's past behaviour and characteristics, such as their age and confidence. In this thesis, I focus on children's attention to two speaker characteristics, both involving their speech properties: accent and filled-pause speech disfluencies (e.g., umm and uhh). Past research suggests children prefer to learn from informants who *speak like them* (Kinzler et al., 2011) and from informants who *speak fluently* (White et al., 2019).

Current Research. In the present experiments, I aim to better understand the strength and malleability of children's use of accent and filled-pause speech disfluencies as cues to informant reliability. In Experiment 1, I ask whether 4-7-year-olds adjust their inferences about disfluencies when they are produced by non-native accented speakers. In Experiment 2, I ask which of the two cues (accent or fluency) 4-6-year-olds consider a more reliable indicator of an informant's knowledge.

Method. In Experiment 1, children were presented with two German accented male speakers. Each speaker first presented familiar objects, either fluently or disfluently. In the test phase, each speaker labelled different novel objects with a novel name (e.g., mido), fluently or disfluently. Children were asked which item was labelled correctly. In Experiment 2, children were presented with one native English speaker and one German accented speaker. As in Experiment 1, one speaker produced fluent utterances and one speaker produced disfluent utterances.

Results. Results from Experiment 1 demonstrate that, overall, children show a weak preference for endorsing fluent speakers. These results suggest that children do not adjust their interpretation of disfluencies based on a speaker's nativeness. Preliminary results from Experiment 2 suggest that children consider a speaker's nativeness to be a stronger cue to their reliability than the speaker's fluency.

Acknowledgements

The two years I have spent working on my thesis have been captivating, intriguing and fascinating, but they have also challenging, demanding, and tiring. Going through this experience has given me opportunities to learn so much about myself and the world and I am so glad I was given the privilege to go through it. It was, however, not without the support and encouragement of the people surrounding me; therefore, I would like to take the time to thank everyone that has been there as I navigated my way through this journey.

To my mother, thank you for showing me the strength and resilience to fight for what I want in this world. You have built me into the strong human being I am today, and I cannot begin to explain how appreciative I am of you.

To my brother, Sergio, thank you for never failing to lighten my mood when I ever felt down about myself or the things going on in my life. You never fail to remind me that life is full of beautiful little pockets of wonderful.

To all the friends and family that have given me a warm and welcoming space to write and work on my thesis. I had the privilege and honour to complete this research and my thesis in multiple environments over the course of the last year and wherever I was and whatever I was going through, there was never a moment I didn't feel loved, supported, and comforted to continue working toward my goals.

To all my fellow graduate students and colleagues, we have been through an interesting year, and it has not been an easy road. Despite the distance, there was never a moment I did not feel connected with you. Thank you for all the motivating chats, the study dates, the advice, the random calls, and most importantly, for all the laughs. I could not imagine having gone through this experience without each and every one of you.

To anyone who helped me during the collection and organization of data, whether that be the organizing of participation, allowing me to go ahead of you in our draft meetings or just putting a smile on my face during the first few minutes of our zoom meetings. A very special thanks to Olivia Halley and Ashley Avarino, who were both rockstars in helping me collect as much data as possible before the submission of this thesis.

To all the amazing children and families that took part in this study, without you, none of this would be possible.

To all of those that let me record your voice for my studies. Finding speakers is one of the hardest parts of the process and you made it just that much easier for me. Thank you.

Finally, last but most definitely not least, I would like to turn the spotlight to my wonderful supervisor, Dr. Katherine White. My time spent working with you has provided me with an experience that has allowed me to grow, not only as a researcher but also as a person. I cannot begin to thank you for allowing me the time, space and trust to discover my interests and for giving me the guidance and the encouragement to explore my strengths, my weaknesses and myself. I hope to continue working with you and contributing to your research in the coming years!

Dedication

I would like to dedicate this thesis to my father. You were taken from this world way too soon. I hope wherever you might be, you can see this and are smiling.

Table of Contents

List of Figures

List of Tables

Introduction

Children are constantly faced with new information and therefore are constantly learning. In fact, their main job in the first few years of life is to learn. Every day, they receive new information about their environment: new people, new words, new objects, and so on. How do children learn this information? Some of it is learned through their own observations, but much of it is learned through what other people show or tell them. The process of learning from other people is, however, not an easy one. Potential information sources have varying levels of knowledge, they make mistakes, they may be misinformed, or they might just be deceptive. Therefore, not all information shared is useful to the learner. Considering this, a crucial question arises: how does a child determine when and if a speaker is sharing valuable, relevant, and accurate information? Research suggests that children use a number of cues to gauge a speaker's reliability. That is, they don't just learn from anyone – instead, they pay attention to a speaker's past behaviour, as well as to other characteristics that may signal their knowledge, in order to determine whether the speaker has valuable information. In other words, they learn selectively from others.

In this thesis, I will first review some previous research on children's use of a speaker's past behaviour and their verbal and non-verbal characteristics as cues to their reliability. I will then present the current research, which aims to investigate two specific speech cues that children have been shown to use when they are gauging an informant's reliability: their production of disfluencies and their accent.

Children's Judgements of Informants based on Past Behaviour

Accuracy. One cue that has been extensively investigated in the selective trust literature is children's sensitivity to a speaker's past accuracy. This research explores whether children are

capable of keeping track of a speaker's past accuracy and whether they use this to help them determine if the speaker might have accurate information for them in a future learning scenario. Past accuracy is a particularly direct cue to a speaker's potential reliability. A speaker who has been consistently wrong in the past is likely to be wrong again in other situations.

Imagine that you are on a hike with two friends. During the hike, friend A has been consistently correct in naming the common plants you have seen, and friend B has been consistently or even just mostly incorrect. Suddenly, you run into a plant you have never seen before. Who are you most likely to turn to, to give you information about this new plant? Friend A, who has been accurate throughout your hike, or friend B, who has not? If you are a good selective learner, then your answer should be friend A, who appears to know a lot more about plants. In this simple example, we see that tracking a speaker's accuracy and using this information for future learning is something that adults would do quite readily. Are children capable of tracking a speaker's past accuracy in this way? Research suggests that starting at as young as 3-years-old children can and do track this type of information (Birch et al., 2008; Corriveau et al., 2008; Koenig & Harris, 2005; Tenney et al., 2011).

In one such study, Birch, Vauthier and Bloom (2008) investigated children's tracking of accurate information by applying a version of a trust paradigm, a paradigm often used in experiments investigating children's choices of informants. They presented 3- and 4-year-old children with two puppets and four familiar objects. After the puppets introduced themselves to the child, they both labelled each familiar object. Puppet A labelled the four familiar objects correctly (e.g., referring to a ball by saying "I think that's a ball. Yeah, that's a ball"), while Puppet B labelled them incorrectly (e.g., referring to a ball by saying "I think that's a book. Yeah, that's a book"). After the familiar object labelling, the experimenter presented four pairs

of novel objects. During the presentation of each object pair, one object appeared in front of Puppet A and the other appeared in front of Puppet B. Each puppet labelled the objects in front of them with the same label (e.g., "I think that's a ferber. Yeah, that's a ferber. Do you see the ferber?"). Afterward, the experimenter closed her eyes, cupped her hands in front of her and asked "Can you give me the ferber? Where's the ferber?". Both 3- and 4-year-old children were significantly more likely to hand over the object labelled by the previously accurate puppet. These results, along with other research (Corriveau et al., 2008; Koenig & Harris, 2005; Tenney et al., 2011) suggest that children use a speaker's past accuracy as a cue to their future credibility.

Expertise. Let's, again, go back to our hiking example. Say you haven't heard any plant names yet from either friend, but you know that friend A is a botanist and friend B is not. In this case, you are most likely going to turn to your botanist friend to give you information about an unknown plant. This same reliance on expertise is also evident in children's learning. That is, if a child learner is aware of a speaker's expertise in a given domain, they take this into account when determining whether they want to learn from them (Aguiar et al., 2012; Lane & Harris, 2015; Lutz & Keil, 2002). For example, 3-, 4-, and 5-year-olds know to consult a doctor and not a mechanic about how to fix a broken arm (Lutz & Keil, 2002).

Lane and Harris (2002) applied a version of the trust paradigm outlined in the study above, but rather than asking children to consult with an informant regarding the labeling of novel objects, children were asked to make judgements about claims experts were making. Children were first introduced to expert informants (e.g., an animal expert and an artifact expert). They were then shown videos of the experts making claims that were either related to their expertise or unrelated to their expertise and were asked to make judgements about which claims

were true. Children were more likely to accept statements made by an expert on the given topic, suggesting that children's trust in informants is dependent on their expertise.

Ignorance. If we think again about how, as adults, we might gauge a speaker's reliability, another cue we might use is whether someone has even been able to offer information in the past at all. Going back to our hiking example, if friend A is consistently labelling common plants correctly, and friend B never offers up any information at all (e.g., saying things like "I don't know what this is"), naturally we are going to turn to friend A when seeking information about an unknown plant. Research shows that this is also a cue children use to gauge a speaker's reliability. For example, they show a preference for a knowledgeable speaker over a speaker who demonstrates ignorance (Koenig & Harris, 2005; Sabbagh & Baldwin, 2001; Sabbagh & Shafman, 2009).

In one such study, Koenig and Harris (2005) introduced 3-, 4-, and 5-year-olds to a knowledgeable speaker, who referred to familiar objects with object labels, and an ignorant speaker, who referred to familiar objects by saying "I don't know what it's called". During the test phase, both informants labelled the same novel object with conflicting names. Children were then asked which speaker labelled the object correctly. Results showed that children had a significant preference for the knowledgeable speaker over the ignorant speaker. These results demonstrate that children treat both inaccurate and ignorant speakers as unlikely to have good information (though they weigh inaccuracy more strongly, Koenig & Harris, 2005).

Confidence. Adult listeners also pay attention to verbal and non-verbal cues that indicate a speaker's confidence in the information they are sharing. For example, the speaker may include explicit statements of certainty, such as "I guess" or "I think", or they may shrug their shoulders as if to indicate that they are unsure. Imagine that friend A names a tree by saying, "This is a

white birch tree" and friend B names the same tree by saying "I think this is a white spruce tree". Who are you more likely to trust? Most likely, the friend who demonstrated confidence in the information they shared (it should be noted that this is not always a good strategy, as some people tend to be overconfident!). Do children make similar judgements about a speaker's reliability based on their confidence? Research suggests that they do (Moore et al., 1989; Tenney et al., 2011). In one study, Jaswal and Malone (2007), demonstrated that 3-year-old children were more likely to accept labels from a speaker who made assertions demonstrating confidence (e.g., "this is a key") than a speaker who made assertions demonstrating a lack of confidence (e.g., "I think this is a key").

Filled-Pause Speech Disfluencies. Another aspect of an individual's speech that can be used as a cue to their confidence is the presence of filled-pause speech disfluencies, such as "um" and uh". Filled-pause speech disfluencies can indicate a lack of certainty or knowledge (Brennan & Williams, 1995) and, in some situations, adult listeners interpret them that way (Brennan & Williams, 1995; Christenfeld, 1995; Amick et al., 2017; Von Tiling, 2011).

Just as children use more overt verbal and non-verbal cues to certainty during learning, they also use speech disfluencies to guide their selection of informants (White et al., 2019). In this study, children were presented with two puppets; one puppet spoke fluently and the other spoke disfluently. During a familiarization phase, the puppets each named three familiar objects (e.g., "This is a clock" vs. "This is a uhhh clock"). In the test phase, each puppet pointed at and labelled one of two novel objects, with the same label (maintaining their fluency from the familiarization phase). Children were then asked which of the two objects was labelled correctly. Results showed that children endorsed the fluent speaker's label for the novel object significantly more often than they endorsed the disfluent speaker's label. This suggests that children make

judgements about an informant's reliability based on speech fluency. That is, when no other information is provided about an informant, children prefer learning from a fluent speaker over a disfluent speaker.

However, disfluency can also arise for non-certainty related reasons. For example, research demonstrates that filled-pause speech disfluencies often occur in particular situations and/or at particular time points in speech, such as when a speaker is retrieving more difficult information or when they are engaged in utterance planning (Bortfeld et al., 2001; Fox Tree & Clark, 1997; Schacter et al., 1991). Indeed, adult listeners are sensitive to these distributional regularities and use them to make predictions about upcoming referents – they expect disfluencies to be followed by referent descriptions that are harder to access or formulate (Arnold et al., 2003; Arnold et al., 2007; Arnold et al., 2004). Surprisingly, children are also sensitive to this property of disfluencies, and make predictions about upcoming linguistic material (Kidd et al., 2011; Owens & Graham, 2016). Children can also adjust their predictions based on information about the knowledge of specific speakers (Orena & White, 2015).

Given that listeners can interpret disfluencies in different ways and that they are, at best, an indirect cue to speaker knowledge, in my thesis I explore the strength and malleability of children's use of filled-pause speech disfluencies as a cue to informant reliability.

Children's Judgements of Informants based on Speaker Characteristics

In addition to using a speaker's past verbal or non-verbal behaviour to determine their reliability, children pay attention to other characteristics of individuals that may serve as indirect cues to their knowledge, such as age or group status.

Age. Age is a speaker characteristic that children pay close attention to when deciding who to trust. Overall, research has shown that children demonstrate a strong preference for adult informants over child informants (Jaswal & Neely, 2006; Rakoczy et al., 2010; Taylor et al., 1991; Wimmer et al., 1988).

Children's assumptions about adult omniscience are demonstrated in a study by Taylor, Cartwright, and Bowden (1991), where 4- and 5-year-old children were presented with pictures of three informants: a baby, a 4-year-old, and a grown-up. Children were then presented with a series of questions and were asked to identify which of the informants would know a given fact. Some of the questions were simple and consisted of facts both the adult and the child, but not the baby, would know (e.g., what a rabbit looks like); other questions were difficult and consisted of facts only the adult would know (e.g., what a lemur looks like). Results demonstrated that 4- and 5-year-old children expected only the adult to be knowledgeable about the difficult facts, but both the adult and the child to be knowledgeable about the simple facts. In a more recent study, Jaswal and Neely (2006) presented children with both an adult informant and a child informant in a novel word learning scenario. When no other information was provided about the informants, children preferred to learn novel object names from the adult. The findings of these two studies, and others (e.g., Rakoczy et al., 2010; Wimmer et al., 1988), demonstrate that children have a belief that adults are generally more knowledgeable and reliable informants than children.

In general, it is reasonable, smart, and safe for children to operate on a belief of adult omniscience. A 4-year-old might reason that "adults are always presenting me with new information and up to this point, I have been able to trust the information they give me, so when faced with something new, it would be most efficient that I continue to seek out information from adults." As adult learners, however, we know that such a rigid criterion for choosing informants isn't always an optimal learning strategy; age is not a direct indicator of someone's knowledge. Indeed, different types of informants may have different types of knowledge, and this includes individuals of varying ages. If you want to know about the history of the differing political parties in your country, who are you going to ask? Probably another adult (and maybe one who is significantly older than you). On the other hand, if you want information regarding the new character, Cap'tn Turbot, on the children's show *Paw Patrol*, would it still be optimal to ask another adult? Or might it make more sense to ask your 5-year-old cousin? If you are a good selective learner, then your answer should be your 5-year-old cousin.

It turns out that children also show this kind of flexibility in learning from informants of different ages. In one study, VanderBroght and Jaswal (2009) presented children with an adult and a child informant. Rather than presenting participants exclusively with topics adults are generally more knowledgeable about, they presented topics that child informants might be more knowledgeable about: toys. Children were significantly more likely to direct toy questions to the child informant than the adult informant (VanderBroght & Jaswal, 2009). These results suggest that children are able to adjust their baseline preference for learning from adult speakers based on the learning situation at hand.

Children are also able to override this general belief of adult omniscience if more information is given about an adult's knowledge state. Jaswal and Neely (2006) investigated this,

when they, again, presented 4- and 5-year-olds with an adult and a child informant. Here, one of the informants demonstrated inaccuracy during the labelling of familiar objects. Participants were then asked to choose an informant to help them learn labels for novel objects. When the adult speaker was inaccurate labelling familiar objects and the child speaker was accurate, children preferred to learn novel object names from the child (Jaswal & Neely, 2006).

These studies demonstrate that, although children tend to have a general preference for adult informants, this preference is not rigid. Children are able to update their preferences for informants in order to potentially maximize learning when either domain or previous behavior suggests that a younger informant may provide better information. More generally, this suggests that children can weigh cues to reliability differently, depending on the situation.

In- vs. Out-Group Status. Another speaker characteristic that listeners pay close attention to when judging informants is their inferred social group. Every day, whether done consciously or not, adults make judgements about individuals based on their social group. We may use group status to determine who to befriend or who to seek information from, under the assumption that members of our in-group are more similar to us or have more relevant information for us (whether this assumption is justified or not). There are many possible indicators of an individual's group status. For example, we may ask whether this person speaks like us or lives close to us. Judgements about whether an individual is a member of our group, in turn, lead us to make inferences regarding the individual's knowledge and experience, and whether they have relevant information for *us*. For adults, a speaker's accent is a particularly strong indicator of their social group, providing key information about their background, such as where they originate from (Derwin & Munro, 2009).

Children show an early sensitivity to accent. Infants as young as 10 months prefer to interact with native accented individuals and 5-year-olds preferred to be friends with individuals who speak in their native accent (Kinzler et al., 2007). These preferences may be because children make inferences about an individual's group status based on their accent and have social preferences for individuals who are part of their in-group.

Children also make similar (though less detailed) inferences as adults about a speaker's geographical background based on their accent (Hirschfeld & Gelman, 1997; Weatherhead et al., 2016, Weatherhead et al., 2018). Results from one study by Weatherhead et al. (2016) showed that 3- to 5-year-old children infer that two speakers who share the same accent live in the same place. They also found evidence that 4- and 5-year-old children infer that speakers who share the same accent might share the same cultural norms and that speakers with different accents might have different cultural norms.

How do these inferences about a speaker's geographical background and group status guide children's learning behaviour? By 3-4-years old, children use an individual's group status to help them determine whether they might be a reliable source of information (Harris & Corriveau, 2011; MacDonald et al., 2013; Zhang & Sylva, 2021). In Kinzler, Corriveau, & Harris (2011), children were presented with two speakers: one native accented speaker and one foreign accented speaker. Both speakers then demonstrated the use of a novel object, in differing ways. They found that children selectively endorsed the non-verbal information from the native accented speaker. These results suggest that children use "nativeness" as a guide in informant selection. In this case, they may have inferred that an individual from the same group would be more likely to use the object in a way that would be relevant for *them*.

It should be acknowledged that it is possible these preferences could be due to a simple preference for familiar speech, rather than an assumption about the quality of information from speakers with different backgrounds. However, research conducted by Begus, Gliga, and Southgate (2016) suggests that preferences for learning from native-accented speakers are indeed due to an inference about information quality. In a first study, Begus and colleagues demonstrated that 11-month-olds' theta activity, as measured by EEG, increases in response to an expectation of information. This pattern of theta activity was then measured in infants' observations of both native-accented and non-native accented informants. Theta activity increased significantly more when observing a native-accented informant when compared to a non-native accented informant. This research suggests that infants show a preference for native-accented informants because they expect to receive useful information from them. Similarly, children likely prefer to learn from native-accented speakers because they believe these speakers will have the most relevant information for *them*. In my thesis, I explore the strength of accent as a cue to informant reliability.

Current Research

In this thesis, I focus on two specific speech characteristics that children use when determining a speaker's reliability: fluency and accent. As discussed, research suggests that children prefer to learn from informants who *speak fluently* and from informants who *speak like them*. However, neither of these cues is a direct indicator of what an individual knows. Both involve inferences – accent may be seen as an indication of knowledge because children make inferences about the individual's background. And in the case of fluency, children may infer that

a lack of knowledge is a reason for a disfluency. Below I summarize the complications inherent in using these two cues:

Disfluency. A speaker's production of disfluencies might not necessarily be due to their lack of knowledge of the topic. Instead, they could be disfluent for other reasons. For example, they may be distracted, or they may be experiencing difficulty with utterance planning or retrieving information (Bortfeld et al., 2001; Fox Tree & Clark, 1997; Schacter et al., 1991). Word retrieval and planning difficulties are processing issues and do not necessarily indicate a lack of knowledge. One very striking example of this dissociation between knowledge and disfluency production is in the speech of second language learners. A non-native speaker is likely to be less proficient in their second language than in their first. They may also produce more disfluencies in their second language. This, however, does not mean they are less knowledgeable. They might just be experiencing greater difficulty in planning their utterances and retrieving words in their second language. In fact, adult listeners seem to realize this. Research by Bosker and colleagues demonstrates that adults adjust their interpretations of disfluencies in some situations when they are produced by non-native speakers (Bosker et al., 2019; Bosker et al., 2013), suggesting that adults are aware that a non-native speaker's use of disfluencies may differ from a native speaker's.

Accent. It is not necessarily the case that a speaker with a different accent has less knowledge about a particular topic. It does not even necessarily indicate that they have less relevant information. For example, it could be that, despite growing up in a different country, they now live in the same country as the person who needs information and are familiar with the conventions of that country. In cases such as this, using accent alone to determine which speaker

is more reliable is not the best strategy. These complications in the interpretation of fluency and accent lead to the two research questions I will focus on in this thesis:

1. Do children adjust their inferences about disfluencies when they are produced by non-native accented speakers? I aim to answer this question in Experiment 1.
2. Children make inferences about a speaker's reliability based on an informant's accent (Begus et al., 2016; Kinzler et al., 2011) and an informant's fluency (White et al., 2019). However, both cues are indirect cues to an informant's knowledge – they could indicate a lack of relevant knowledge, but they may not. In Experiment 2, I ask which of the two cues children weigh more heavily during learning. In other words, which cue do children consider a more reliable indicator of knowledge, when both cues are present?

Experiment 1

In Experiment 1, I aim to investigate how children treat disfluencies when they are produced by non-native speakers. Are children flexible in their interpretation of disfluencies as cues to reliability? Research shows that children's use of disfluencies is flexible in one direction - they adjust their processing of a speaker's disfluencies when given direct evidence that the speaker lacks knowledge (Orena & White, 2015). But do children make adjustments in the other direction? Do they adjust their disfluency-based inferences about a speaker's knowledge when there is an (indirect) alternative reason for the speaker's increased use of disfluencies, such as the fact that they are a non-native speaker? There are at least two possible reasons why children may not engage in this adjustment. One is if children consider disfluency to be such a strong cue to a speaker's confidence/knowledge that they rely on it even in situations when a speaker might be disfluent for other reasons. A second reason that children might not alter their use of disfluencies for non-native speakers is if they are not yet aware that non-native speakers can be more disfluent because of processing issues. The predictions for Experiment 1 are as follows:

1. If children are aware that a non-native speaker's use of disfluencies might differ from that of native speakers (that retrieval/planning problems are more likely), then disfluencies should be ignored when they choose between two non-native informants. In this case, children's choices between fluent and disfluent non-native informants will be at chance.
2. If children are not aware that a non-native speaker's use of disfluencies might differ from that of native speakers or believe that disfluency is a particularly strong cue to reliability, then the pattern of results should not differ from past research on the use of disfluency during selective trust tasks. In this case, children should choose the fluent speaker at above chance levels.

Methods

Participants. Sixty-nine children participated (mean age = 68 months): 13 4-year-olds (7 females, *M* = 54 months, range = 48-59 months), 34 5-year-olds (16 females, *M* = 65 months, range = 60-71 months), 17 6-year-olds (8 females, *M* = 78 months, range = 72-83 months) and 5 7-year-olds (3 females, *M* = 91 months, range = 86-95 months). The discrepancy in sample sizes across the age range is due to challenges in participant recruitment. Data collection is ongoing in the age groups with smaller sample sizes. Children participated over Zoom through the University of Waterloo's CORAL (Child Online Research Activities Lab) program. All parents reported English as their child's only or strongest language. Participants had no known auditory or visual conditions; their parents gave informed consent for their participation and were entered into a $50 Amazon gift card draw as a token of appreciation for their participation. An additional 18 participants were tested (4 4-year-olds, 9 5-year-olds, 4 6-year-olds, and 1 7-year-old), but their data were not included in the analysis due to changes in the testing procedure or high distraction during the Zoom call.

Materials

Audio Stimuli. Two male bilingual speakers of English and German recorded stimuli with a German accent. Each speaker recorded two introduction sentences, twelve fluent familiarization sentences, twelve disfluent familiarization sentences, two fluent test sentences, and two disfluent test sentences (see Table 1 for a full list of sentences).

Visual Stimuli. Visual stimuli consisted of four images of male actors: two actors for trial 1 and two actors for trial 2. Actors within each trial were matched in appearance. Visual object stimuli included twelve familiar objects (cookie, bird, chair, clock, pig, flower, door, lion, book, balloon, cat, key) and four novel object images retrieved from Horst and Hout's (2014) *Novel*

Object and Unusual Name (NOUN) Database 2nd edition (see appendix D for images of visual stimuli).

Table 1.

List of sentences and familiar and novel object names for Experiments 1 and 2.

Introduction	*Hello, let's play a game* *Hi there, yes, let's play*		
	Carrier sentence		**Target word**
	E1	**E2**	
Familiarization (fluent)	*That is a*	*That is a*	*cookie, bird, chair, clock, pig, flower, door, lion, book, balloon, cat, key*
Familiarization (disfluent)	*That is a* ***umm***	*That is a* ***uhh***	
Test (fluent)	*This is a*	*This is a*	*mense, wilp*
Test (disfluent)	*This is a*	*This is a* ***uhh***	

Accent ratings. To assess the strength of the speakers' foreign accents in Experiment 1 and 2, ratings were collected from a sample of adults (n=20)[1], who completed an online survey rating each speaker's accent in relation to a standard Canadian English accent. Participants were

[1] To examine whether children were sensitive to the foreign accents, we also got ratings from a small sample of child participants (n=12, mean age = 65 months). A larger sample was intended, but not collected due to a slowdown in recruitment. Children were shown a single trial with images of two cartoon boys next to each other on a PowerPoint slide (see Appendix B for procedure). They heard two speakers, each producing the passage heard by the adult raters: "Once upon a time there was a rabbit that lived in the forest. He was a very hungry rabbit and he loved to eat carrots." One of the speakers had a native Canadian accent, the other was either speaker 1 or speaker 2 from Experiment 1. After hearing both speakers, children were asked "which boy speaks more like you". A score of 1 was assigned if children chose the Canadian speaker and a score of 0 was assigned if they chose the non-native speaker. Even in this small sample, children were significantly more likely to choose the native Canadian speaker (mean choice = .83, $t(11) = 2.97$, $p = .013$), indicating that they could discern the difference between the native and non-native accents. However, given the small sample, we were unable to see whether children also had an easier time detecting the foreign accent for one of the speakers.

asked to rate four speakers' accents (1 native accent and 3 German accents). These were the two speakers from Experiment 1 (both German accented) and the two speakers from Experiment 2 (one native and one German accented). A single passage was presented from each speaker: "Once upon a time there was a rabbit that lived in the forest. He was a very hungry rabbit and he loved to eat carrots." Speakers were presented in random order. Participants were asked to give a speaker a score of 1 if the speaker's accent was similar to a standard Ontario/Canadian accent and a score of 5 if the speaker's accent was very different from a standard Ontario/Canadian accent. On a scale of 1-5, the two speakers in Experiment 1 received ratings of 4.45 and 3.3, respectively, indicating that one of the speakers had a stronger foreign accent. Both of the speakers received ratings that differed significantly from the native speaker's mean rating of 1.65 ($t(19) = -13.61$, $p < 0.001$; $t(19) = -8.43$, $p < 0.001$). The German accented speaker in Experiment 2 received a rating of 4.29, which also differed significantly from the native speaker ($t(19) = -12.70$, $p < 0.001$).

Design and Procedure

Children were tested by researchers from the University of Waterloo's CORAL team. Participants joined a Zoom meeting and at the beginning of each session, parents and children were led through a set up and calibration procedure to determine what device was being used and to ensure that stimuli appeared correctly on the participants' screen (e.g., full screen mode, minimizing researchers' video boxes, etc.). Before starting any activities with children, the researcher asked parents not to help their children answer any of the questions asked during the session (e.g., by putting themselves in a position where they could not see the screen) and reminded children that if they wanted to stop playing to let the researcher know.

The study was presented over fourteen Microsoft PowerPoint slides. First, participants went through an audio and visual check to ensure they could hear and see the stimuli. Following this check, two experimental trials were presented in succession. The same voices were used in both trials, however, different actor images appeared with the voices on trial 2. Within each trial, participants went through a novel object familiarization phase, a speaker familiarization phase, and a test phase. At the end of the two trials, participants were asked an explicit judgment question.

Audio and visual check. Participants were presented with the audio from a female, native English speaker who asked children to tell them "which box the flower is in" (the blue box on the upper left corner of the screen) and "which box the tree is in" (the orange box on the upper right corner of the screen), see Appendix A and C for a detailed outline of the procedure and script. If children were able to answer these questions correctly, it was determined that (1) they could hear the audio clearly, (2) stimuli were being presented correctly on their screen, and (3) they had no trouble naming the respective colours of the boxes (which were used for the remainder of the experiment).

Novel object familiarization phase. At the start of each experimental trial, children were familiarized with two novel objects. Here the experimenter provided the following instructions: "Here are two new toys. We have to figure out what one of these toys is called.". The side of the novel objects was counterbalanced across participants. See Appendix A and C for a detailed outline of the procedure and script.

Speaker familiarization phase. Following the novel object familiarization, two male, German accented speakers were introduced with actor images. The experimenter introduced the speakers by saying: "Here are two guys. We're going to listen to them tell us about some things,

and then you can pick one of them to help us figure out the name of the new toys!". Each speaker then introduced themselves in succession (first speaker: "Hello, let's play a game", second speaker: "Hi there, yes, let's play"). After these introductions, each speaker presented different sets of three familiar objects. One speaker presented the objects fluently (e.g., "That is a cookie", "That is a bird", "That is a chair") and one speaker presented the objects disfluently (e.g., "That is a umm clock", "That is a umm pig" "That is a umm flower"). The same voices were used in both trials and voice fluency remained consistent for the two trials. The set of familiar objects presented first, which familiar objects each speaker presented, and which voice was speaking fluently versus disfluently were all counterbalanced across participants.

Test phase. After each familiarization phase, children went through a test phase. Here, speakers each presented one of the novel objects from the start of the trial. The experimenter gave instructions by saying: "Remember these toys? One of these is a MENSE! Let's listen to what the guys say!". Each speaker then labelled the object below them with the same label, either fluently (e.g., "This is a mense") or disfluently (e.g., "This is a umm mense"). Speaker fluency remained consistent from the familiarization trials. After the speakers labelled the objects, the experimenter asked the child: "Which one is a MENSE? The one in the blue box or the one in the orange box?". The side of the novel objects, the novel object labels, the trial order and the side of speakers was counterbalanced across children. Speakers remained on the same side throughout the trial.

Explicit judgement question. After the second trial, participants were asked an explicit judgement question: "Was there anything funny about the way the guys were talking?". This was to determine whether children explicitly noticed non-native speakers' disfluencies. If children made note of the speaker disfluency, they were further asked: "Why do you think they were

saying umm?" or "Why do you think they were talking like that?". Past research has suggested that children believe speakers use disfluencies because they are less knowledgeable (White et al., 2019). We wanted to determine whether this would hold when listening to disfluencies produced by non-native speakers. This question was asked only after the second trial, to avoid explicitly cuing children to speaker fluency, which might have influenced their speaker choices on the second trial.

Language and Accent Exposure. Parents provided information about their child's language and accent exposure either within the Zoom call after the completion of the study or through filling out a questionnaire sent in an email after participation.

Results

Full sample of 4- to 7-year-olds

Before examining children's choices of fluent versus disfluent speakers, I sought to determine whether there was a preference for one speaker over the other (given the difference in accent strengths). Each child received a score of 0 when they endorsed information given by Speaker 1 and a score of 1 when they endorsed information given by Speaker 2. Responses from the two trials were added, giving each participant a total score out of 2. Overall, children did not show a significant preference for either speaker ($M = 1.03$, $t(68) = 0.314$, $p = 0.754$).

To calculate participants' choices in endorsing information from fluent or disfluent speakers, each child received a score of 0 when the information given by the *disfluent* speaker was endorsed and a score of 1 when the information given by the *fluent* speaker was endorsed. Responses from the two trials were added, giving each participant a total score out of 2. Overall, children were more likely to choose the fluent speaker than the disfluent speaker (with a mean

score of 1.16 out of 2), see Figure 1. However, a one-sample t-test against chance showed that this difference did not reach significance ($t(68) = 1.789$, $p = 0.078$).

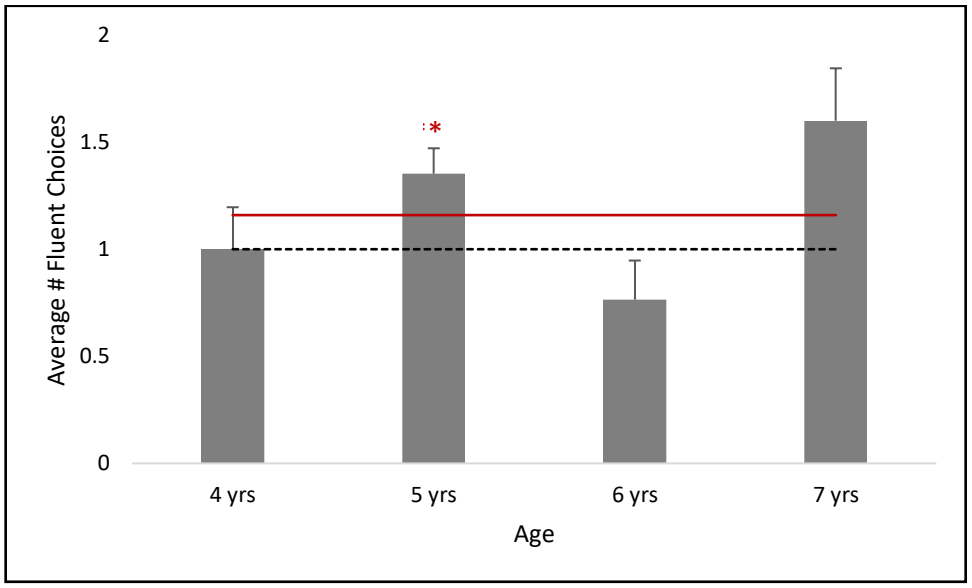

Figure 1. Bar graph representing data for all age groups in Experiment 1. The dotted black line represents chance (chance = 1) and the red line represents the overall average ($M = 1.16$, $SE = 0.09$). Individual bars represent each age group. Error bars represent standard error.

Responses to explicit judgement question. Next, I considered children's endorsements based on their responses to the explicit judgement question ("was there anything funny about the way the guys were talking?") asked at the end of the second trial. White et al. (2019) found that children who explicitly noted the disfluency chose the fluent speaker at above chance levels, whereas those who did not comment on the disfluency chose at chance. In the current study, out of the 69 participants, 24 participants explicitly pointed out the disfluency and 45 participants did not comment on the disfluency. One sample t-tests against chance revealed that children who explicitly pointed out the disfluency had a significant preference for the fluent speaker ($M = 1.29$, $t(23) = 2.29$, $p = 0.032$), but there was no significant preference for children who did not comment on the disfluency ($M = 1.09$, $t(44) = 0.75$, $p = 0.456$). This replicates the previous

finding that children who have more explicit awareness of the disfluency are more likely to use it to assess reliability. However, an independent samples t-test revealed that the two groups did not significantly differ in their preference for the fluent informant ($t(67) = 1.09$, $p = 0.282$).

Influence of language background. Information regarding participants' language and accent exposure was provided for 63 out of 69 participants. For these children, children hearing a language other than English more than 20% of the time were considered multilingual (ranges between 0%-30% are common in child language research). Out of the 63 participants that we received language information about, 22 were multilingual. One sample t-tests revealed that these children had a marginal preference for the fluent speaker ($M = 1.27$, $t(21) = 2.03$, $p = 0.056$), whereas the monolingual children performed at chance in their choice of informant ($M = 1.15$, $t(40) = 1.23$, $p = .225$). However, an independent samples t-test revealed that the two groups did not significantly differ in their choice of informant ($t(61) = .67$, $p = 0.508$).

Influence of accent background. Next, the data was analyzed based on accent exposure (high versus low). Children who heard accents other than the local Ontario accent more than 20% of the time were considered to have high accent exposure. Out of the 60 participants that we received accent exposure information for, only 11 had high accent exposure and 49 had low accent exposure. One sample t-tests revealed that the small sample of children with high accent exposure performed at chance in their choice of informant ($M = 1.09$, $t(10) = 0.36$, $p = 0.724$), whereas children with low accent exposure had a significant preference for the fluent speaker ($M = 1.22$, $t(48) = 2.2$, $p = .033$). An independent samples t-test revealed that the two groups did not significantly differ from each other in their choices of informant ($t(58) = -.54$, $p = 0.589$).

Breakdown by Age

Because the 5- and 6-year-old age groups had more data than the younger and older children, I conducted some further analyses of only these age groups separately. I first asked whether these children had a significant preference for one of the speakers over the other (a score of 1 represents chance). As in the full sample, neither 5- (M = 1.03) or 6-year-olds (M = 0.94) had a significant preference for one of the speakers ($t(33) = 0.22$, $p = .831$; $t(16) = -0.29$, $p = .773$).

I then compared 5-and 6-year-old children's choices based on fluency. Results showed 5-year-old children had a significant preference for the fluent speaker over the disfluent speaker ($M = 1.32$, $t(33) = 2.76$, $p = .009$), whereas 6-year-olds performed at chance in their choice of informant ($M = 0.82$, $t(16) = -0.90$, $p = .382$).

Responses to explicit judgement question. As in the full sample, I then compared whether there were differences in children's choices based on their responses to the explicit judgement question. This was done only for 5-year-olds, as very few 6-year-olds mentioned the speakers' fluency. Out of 34 5-year-olds, 13 participants explicitly pointed out the disfluency and 21 participants did not. One-sample t-tests revealed that the children who explicitly pointed out the disfluency had a significant preference for the fluent speaker ($M = 1.46$, $t(12) = 3.21$, $p = 0.008$), whereas the children who did not point out the disfluency performed at chance in their choice of informant ($M = 1.24$, $t(20) = 1.42$, $p = .171$). An independent samples t-test revealed that the two groups did not significantly differ in their choices of informant ($t(32) = 0.92$, $p = 0.363$).

Influence of language background. Finally, I considered participants' language background. I again did this only for 5-year-olds, because very few 6-year-olds were reported to be multilingual. Out of 34 5-year-olds, 10 participants were reported to be multilingual, and 24

participants monolingual. One-sample t-tests revealed that both multilingual and monolingual 5-year-olds showed a preference for the fluent speaker over the disfluent speaker, though it was only significant for multilingual children (M = 1.50, $t(9)$ = 3.00, p = 0.015; M = 1.25, $t(23)$ = 1.66, p = 0.110). An independent samples t-test revealed that the two groups did not significantly differ in their choices of informant ($t(32)$ = 0.97, p = 0.339).

No analyses were run based on accent exposure as only 6 5-year-old children were reported to have high accent exposure.

Discussion

Results from Experiment 1 demonstrate that there was a trend for children overall, as well as a significant preference for the 5-year-olds alone, to selectively endorse the fluent speaker. This parallels previous findings with native speakers (White et al., 2019) and suggests that children do not adjust their use of disfluency based on a speaker's nativeness. One possible explanation for this is that children do not expect non-native speakers to have different patterns of disfluency than native speakers. This would contrast with other inferences that children make about non-native speakers, such as inferences about their geographical background and that they might have less relevant information than native speakers (Begus et al., 2016; Kinzler et al., 2011; Weatherhead et al., 2016, Weatherhead et al., 2018). An alternative possibility is that children rely so heavily on disfluency as a cue to a speaker's knowledge state that it overrides any assumptions based on accent. In Experiment 2, I wanted to test children's judgements of disfluent and non-native speakers further, by directly pitting fluency and nativeness against each other.

Experiment 2

Children use both accent and speech fluency during learning - they prefer not to learn from speakers who present with either non-native or disfluent speech characteristics. However, accent is not always an indicator of a speaker's knowledge state, and neither is a speaker's fluency. A speaker with a different accent could be as knowledgeable as a speaker with a native accent and a speaker's use of disfluencies could be caused by many reasons other than a lack of knowledge. In Experiment 2, I aim to investigate which of the two speech cues children consider a stronger cue to a speaker's knowledge (or lack thereof) by pitting fluency against accent. Participants heard either fluent native speakers vs. disfluent non-native speakers or disfluent native speakers vs. fluent non-native speakers (between subjects manipulation). The predictions for Experiment 2 are as follows:

1. If children rely more strongly on fluency, they will choose a fluent informant over a disfluent informant even when the fluent informant is a non-native speaker and the disfluent informant is a native speaker. This outcome would suggest that children consider fluency to be a more reliable cue to what someone knows.
2. If children rely more strongly on accent, they will choose a native accented informant over a non-native accented informant, even when the native accented informant is disfluent. This outcome would suggest that children consider accent to be a more reliable cue to what someone knows.

Methods

Participants. Eighteen children participated in the current study: nine 4-year-olds (5 females, *M* = 50 months, range = 45–57 months), six 5-year-olds (6 females, *M* = 62 months, range = 60–68 months), and three 6-year-olds (2 females, *M* = 77 months, range = 72–81 months). As in Experiment 1, all children participated over Zoom through the University of Waterloo's CORAL (Child Online Research Activities Lab) program. Time constraints imposed by the COVID-shutdown mean that very little data have been collected in this experiment so far. We are currently in the process of collecting a sample size similar to that of experiment 1 (*N* = 32 for each age group). All parents reported English as their child's only or strongest language. Participants had no known auditory or visual conditions; their parents gave informed consent for their participation and were entered into a $50 Amazon gift card draw as a token of appreciation for their participation. One additional participant was tested (one 4-year-old), but their data was not included in the analysis due to high distraction during the Zoom call.

Materials

Experiments 1 and 2 were similar in materials and procedure, with some minor changes.

Audio Stimuli. Two male speakers recorded stimuli in English. In contrast to Experiment 1, in which both speakers were non-native German-accented speakers, in Experiment 2, one was a native English speaker, and one was a bilingual German accented speaker (the German accented speaker was not the same as the Experiment 1 speakers). Speakers recorded the same introduction, fluent and disfluent familiarization sentences, and fluent and disfluent test sentences as in Experiment 1 (see Table 1 for a full list of sentences). One additional change that was made between Experiments 1 and 2 was the type of disfluencies children heard. In Experiment 2, speakers produced "uhh" in their disfluent sentences, rather than "umm". This

change was made to be more consistent with past research investigating disfluencies in children (White et al., 2019).

Visual Stimuli. Visual stimuli were identical to what was presented in Experiment 1. The only change made to the visual presentation was the use of more animations within the PowerPoint slide, highlighting which boy was speaking. This was done to ensure that children were able to keep track of which pictures corresponded with which voice, a change made as we refined our online testing procedures (which our lab implemented as a result of the COVID-19 pandemic).

Language and Accent Exposure. As in Experiment 1, parents provided information about their child's language and accent exposure within the Zoom call during participation or through filling out a questionnaire sent in an email after participation.

Design and Procedure

Children were assigned either to a condition in which they heard fluent native speakers vs. disfluent non-native speakers on both trials or a condition in which they heard disfluent native speakers vs. fluent non-native speakers on both trials. The procedure of Experiment 2 was similar to Experiment 1. The only changes made were to the script. To ensure that children were following along, we added some additional instructions to direct children's attention to the corresponding image on the screen (see Appendix A and C for an outline of E2's procedure and script).

Results and Discussion

Results presented are extremely preliminary (N = 18), due to disruption caused by the move to online testing. Data collection is ongoing in all of the age groups.

Fluent versus disfluent speakers. To calculate participants' choices for endorsing information from fluent or disfluent speakers, each child received a score of 0 when the information given by the *disfluent* speaker was endorsed and a score of 1 when the information given by the *fluent* speaker was endorsed. Responses from the two trials were added, giving each participant a total score out of 2. Overall, when both conditions are included (ignoring nativeness), children are not showing a preference for either the fluent or disfluent informants (M = .83, $t(17) = -0.83$, p = .421).

Native versus non-native speakers. To calculate participants' choices in endorsing information from the native or the non-native speaker, each child received a score of 0 when the information given by the *non-native* speaker was endorsed and a score of 1 when the information given by the *native* speaker was endorsed. Responses from the two trials were added, giving each participant a total score out of 2. Overall, when both conditions are included, children are more likely to choose the native accented speaker than the non-native accented speaker (M = .83, $t(17)= 2.68$, p = .016). Although extremely preliminary, the results so far suggest that children consider a speaker's nativeness to be a stronger cue about their reliability than the speaker's fluency.

In this experiment, half of the participants are being asked to choose between a speaker who is both fluent/has a native accent vs. a speaker who contrasts along both dimensions (who is both disfluent/has a non-native accent). In this case, one speaker exhibits two cues consistent with reliability (and the other two cues consistent with unreliability). The other half of the participants are presented with only one cue per speaker – they hear a speaker who is disfluent but native and a speaker who is non-native but fluent. In this case, the two cues confliect. Currently, the data across both conditions together demonstrate that children have a preference

for nativeness over fluency. However, once the full set of data are collected, we will also be able to examine children's preference in each condition separately. If children are using both nativeness and fluency to determine reliability, then children who are in the first condition should show a stronger preference for the native speaker (who is also fluent) than children in the second condition (when the native speaker is disfluent). The results of the second condition alone will tell us whether nativeness is prioritized when the two cues directly conflict.

General Discussion

Research has demonstrated that by the age of 3 years, children are good at keeping track of direct and indirect cues to an informant's reliability (Birch et al., 2008; Koenig & Harris, 2005; Lane & Harris, 2002). One thing children pay close attention to is an informant's speech characteristics, more specifically, speech characteristics that might give them information about the informant's confidence or group status (Moore et al., 1989; Tenney et al., 2011). Speech fluency gives children information about a speaker's confidence (White et al., 2019) and a speaker's accent is a sign of their group status (Kinzler et al., 2007). Children appear to make inferences about speakers using these speech characteristics, preferring to learn from informants who *speak fluently (*presumably because they assume they are more sure of their information) and from informants who *speak like them* (presumably because they assume they are likely to have more relevant information).

Although fluency and accent can provide listeners with valuable information regarding an informant's knowledge state, neither of these cues provides a direct indication of what an individual knows. Use of both a speaker's accent and a speaker's fluency involve *inferences* regarding a speaker's knowledge state. Therefore, these cues should not be used rigidly. The current thesis examines how children weigh these cues during learning.

Experiment 1

Research suggests that adults are aware that a non-native speaker's use of disfluencies may differ from a native speaker's, and, as a result, they adjust their processing of disfluencies produced by non-native speakers (Bosker et al., 2019; Bosker et al., 2013). In Experiment 1, I asked whether children also adjust their use of disfluencies when they are produced by non-native speakers. The results of Experiment 1 suggest that children may not make such

adjustments. Overall, the analyses I ran suggest that children might be treating disfluencies produced by non-native accented speakers the same way they treat disfluencies from native accented speakers.

Why might children be treating non-native speakers the same way they do native speakers when it comes to disfluencies? One possibility is that the strength of the accents used might not be strong enough for children to recognize them as non-native speakers. However, the adult ratings demonstrate that the speakers did have perceptible foreign accents, and the preliminary results from the "Who speaks like me?" task with children suggests that they, too, were sensitive to the difference between the native and foreign accents. I will continue collecting data for this measure to determine whether children were, in fact, sensitive to the non-native accents for both speakers in Experiment 1. A second possibility is that children's judgements aren't affected by the accent because the presence of disfluencies might be such a strong cue to a speaker's knowledge state that children cannot ignore it. Only one previous study has examined whether children use disfluencies as a cue to reliability (White et al., 2019), so the strength of this cue is unknown. Finally, it is possible that children simply do not have enough experience interacting with non-native speakers to realize that they can be more disfluent than native speakers (and that this is not necessarily a reflection of their knowledge). That is, children might need to experience more situations in which they interact with disfluent non-native speakers who demonstrate knowledge despite producing disfluencies. If so, this would mean that children require more experience to make this type of inference about non-native speakers (that their disfluencies do not necessarily reflect knowledge) than they need to make other judgements about accented speakers. For example, even children with little exposure to non-native accents

make inferences about their geographical background (Weatherhead et al., 2016). This will be discussed further in the *high versus low exposure to accents* section below.

Though we found a preference for the fluent speaker in Experiment 1, it should be noted that the effect was non-significant and weaker than was observed in work with native speakers (White et al., 2019). There are a few possibilities that could explain why the effect is weaker. One is that it indicates some sensitivity to possible differences in the use of disfluencies for native and non-native speakers. Although there is an overall (non-significant) preference for the fluent speaker, it could be that some children do alter their use of disfluencies for non-native speakers. In fact, children whose parents reported high accent exposure did not show a preference for either the fluent or disfluent speaker, whereas children whose parents reported low accent exposure did show a preference for the fluent speaker. This is consistent with the possibility that more experience with accents causes children to shift their judgements of disfluencies (and to no longer assume that disfluencies are an indication of a speaker's lack of knowledge). Children with low accent exposure may not yet have shifted their interpretations. This potential effect of experience is very intriguing – as it would demonstrate that children really do reason about the source of disfluencies and are not simply disrupted in their use of disfluencies due to the unfamiliar accent. If the weaker effect of disfluencies were due to disruption from the accent, we would expect children with high accent exposure to be less disrupted and that this group, therefore, would be the ones to treat disfluencies more similarly for native and non-native speakers. It is also interesting to observe that multilingual children did not pattern like children with high accent exposure. Although multilingual children are hearing multiple languages being spoken, this does not mean that they also have high exposure to non-native speakers (some may, but some may be hearing their multiple language input from only

native speakers of each language). Therefore, multilingual children, who may not have had extensive exposure to non-native speakers, may act in a similar way as monolingual children.

The somewhat weaker effect observed here may also be due to the methodological and demographic differences between studies. A major difference between the White et al. (2019) study and the present study is that the data for the present study was collected online. Because this method of testing is novel to our lab, it is likely that we did not have a fully optimal procedure in place when testing started (and, indeed, we made some procedural changes for Experiment 2). Additionally, data collected in online environments has been reported to be noisier than that collected in person (Scott et al., 2017). We also likely tapped a somewhat different population of children in our online testing, even though participants were still recruited primarily from the Kitchener-Waterloo area. And finally, the participants in the present experiments are older, on average, than those tested in our previous lab-based studies of disfluency. This was due to the difficulty in coming up with paradigms to engage younger children online and to our efforts to maximize recruitment (regardless of age). It is possible that disfluency-based learning preferences present less with age (though that remains to be seen once the full sample is collected across age groups).

Experiment 2

In Experiment 2, I wanted to further test the relationship between disfluencies and speaker "nativeness", by pitting the two cues against each other. Children trust a fluent speaker over a disfluent speaker (White et al., 2019), and they trust a native speaker over a non-native speaker (Kinzler et al., 2011) – if the two are in conflict, what happens?

The results of Experiment 2 are extremely preliminary. Children are currently demonstrating a preference for nativeness (collapsed over fluency). If this holds when I examine

the second condition alone (when nativeness and fluency directly conflict), this would suggest that they rely more heavily on nativeness. Below I outline the possible outcomes of the study.

Fluency as the stronger cue

Both accent and fluency can be seen as cues to a speaker's knowledge state. However, they are not direct cues to what a speaker does or does not know. A non-native speaker may indeed be knowledgeable about a particular topic, and, in the case of disfluency, there are other possible reasons for a speaker's use of them (e.g., utterance planning, word retrieval, distraction). Sobel and Finiasz (2020) recently conducted a meta-analysis in which they broke down cues to an informant's reliability in terms of epistemic and non-epistemic cues. Because of the indirect nature of nativeness and fluency, we can place both in the category of indirect "epistemic cues". They *can* tell you something about a speaker's knowledge state, but they could be telling you something else about the speaker as well. Sobel and Finiasz (2020) report that epistemic cues are weighed more heavily by children than non-epistemic cues. However, both accent and fluency are indirect epistemic cues, so it is unclear which will (or should) be prioritized when they are in conflict. If children are rational in their weighting of these cues, then it will be based on some assessment of their relative reliability. For example, it could be argued that a speaker's fluency is a slightly more direct and reliable cue to their knowledge, as it reflects some aspect of their mental state in the moment. In this case, we could predict that children will ultimately prioritize a speaker's fluency over their nativeness. However, given the paucity of data examining the fluency bias in previous research, it is not clear just how strong a cue this is for children.

Nativeness as the stronger cue

Research suggests that children are sensitive to accent differences from an early age (Kinzler et al., 2007) and that they also begin making social judgements about speakers with

different accents, such as inferences about their geographical background (Weatherhead et al., 2016) and their reliability (Harris & Corriveau, 2011; MacDonald et al., 2013; Zhang & Sylva, 2021). Indeed, these findings may be linked – children may choose native speakers to learn from because they assume that they are part of the child's in-group (and live in the same location; Weatherhead et al., 2018). Therefore, they assume they have the most relevant information for them (Begus et al., 2016). In this sense, children may indeed view accent as a very strong indicator of a speaker's knowledge (more so than fluency).

In addition to this, it should be noted that accent is a much more salient and enduring perceptual cue than is fluency. That is, a listener can hear that a speaker is a non-native speaker throughout their utterance, whereas disfluencies only occur at certain points in speech. Children might find it easier to keep track of a more salient cue than a transient cue. This may be an additional reason why the preliminary results indicate that they are weighing nativeness more strongly.

Limitations

One limitation of the experiments in this thesis is how much our paradigm differed from past experiments investigating children's judgements of an informant's reliability based on fluency and accent. Most of these changes were introduced to accommodate the online setting in which we have been collecting our data over the past year due to the lockdowns and restrictions caused by the spread of COVID-19. One such difference is the use of still frame images to represent each speaker, rather than videos of speakers. This was done to accommodate the online environment, where bad internet connections can cause videos to lag. In addition, in a live experiment, the experimenter is able to point to each speaker as they are talking. This is obviously not possible in an online environment and therefore we had to add animations (such as

the box around the speaker) to attempt to draw children's attention to the featured speaker. It could be argued that children may not be able to keep track of which image belongs to which voice. As noted earlier, although children overall are showing a preference for the fluent speaker in Experiment 1, this preference is weaker than has been observed previously, and difficulty linking the voice and image could have contributed to this. Another difference is that children are in an entirely different environment during testing than they have been in the past. Past research has been completed within a lab setting with a controlled environment and few distractions. In contrast, children participated in the current experiments in an online environment from home. Despite our efforts, being in a home setting results in greater variability in distraction levels, which may be affecting the results (especially since these experiments rely on children attending to subtle speech characteristics, rather than, e.g., salient aspects of the visual display).

Another limitation lies within the differences between Experiments 1 and 2. Experiment 1 was designed in the first stages of our lab's shift from in person studies to online Zoom testing. At the time, there were very few language studies that had been done in an online setting, meaning we had little previous work to consult for study procedures that would work in this setting. In fact, we realized over the course of our first term of testing (fall 2020) that some of our procedures would need to be altered to work in the online environment. Our first attempts at data collection on a different disfluency project yielded uninterpretable data (children made the same choice on almost every single trial). We therefore abandoned that project and set up Experiment 1 of my thesis. This procedure proved more suitable. We continued to refine our procedure when moving from Experiment 1 to Experiment 2. These changes included (1) using more PowerPoint animations and making minor changes to the script to ensure children are keeping track of which pictures correspond with which speakers and (2) removing the labeling of

box colours in the test phase (e.g., Exp 1: “Which one is the WILP? The one in the blue box or the one in the orange box?”; Exp 2: “Which one is the WILP?”), to ensure that the experimenter’s intonation when labelling the box colour would not influence the participant’s responses.

Conclusion

Children use multiple cues to determine whether an informant is trustworthy, including speech characteristics that cue a speaker’s confidence and their group status. Children prefer to learn from informants who speak fluently and informants who speak like them. The current research aimed to contribute to this literature by examining the strength and malleability of these speech cues.

Results from Experiment 1 suggest that 4- to 7-year-old children do not treat disfluencies differently when they are being produced by non-native speakers, at least not in a word learning scenario. However, additional analyses suggest that children may begin to treat disfluencies differently from non-native speakers if they are exposed to accents on a regular basis. This suggests that children may indeed adjust their use of disfluencies as cues to speaker reliability once they have built up knowledge that non-native speakers use disfluencies differently. Preliminary results of Experiment 2 suggest that 4- to 6-year-old children weigh a speaker’s nativeness more strongly as a cue to their reliability than the speaker’s fluency. Future research should further explore the strength of children’s reliance on these two speech cues, as well as how this is influenced by children’s different language experiences.

References

Aguiar, N. R., Stoness, C. J., & Taylor, M. (2012). The development of children's ability to fill the gaps in their knowledge by consulting experts. *Child Development, 83*, 1368–1381.

Amick, L. J., Chang, S.-E., Wade, J., & McAuley, J. D. (2017). Social and cognitive impressions of adults who do and do not stutter based on listeners' perceptions of read-speech samples. *Frontiers in Psychology, 8*, 1148.

Arnold, J. E., Fagnano, M., & Tanenhaus, M. K. (2003). Disfluencies signal theee, um, new information. *Journal of Psycholinguistic Research, 32*, 25–36.

Arnold, J. E., Hudson Kam, C. L., & Tanenhaus, M. K. (2007). If you say thee uh you are describing something hard: The on-line attribution of disfluency during reference comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*, 914–930.

Arnold, J. E., Tanenhaus, K. K., Altmann, R. J., & Fagnano, M. (2004). The old and thee, uh, new: Disfluency and reference resolution. *Psychological Science, 15*, 578–582.

Begus, K., Gliga, T., and Southgate, V. (2016). Infants' preferences for native speakers are associated with an expectation of information. *PNAS, 133(44)*, 12397-12402.

Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition, 107*, 1018–1034.

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),* 123-147.

Bosker, H.R., van Os, M., Does, R., and van Bergen, G. (2019). Counting

'uhm's: how tracking the distribution of native and non- native disfluencies influences online language comprehension, *Journal of Memory and Language*, 106, 189-202.

Bosker, H. R., Quené, H., Sanders, T., & De Jong, N. H. (2013). The processing of disfluencies in native and non-native speech. *Talk presented at The 19th Annual Conference on Architectures and Mechanisms for Language Processing (AMLaP 2013). Marseille, France. 2013-09-02 - 2013-09-04*.

Brennan, S. E., & Williams, M. (1995). The feeling of another's knowing: Prosody and filled pauses as cues to listeners about the metacognitive states of speakers. *Journal of Memory and Language*, 34, 383–398.

Christenfeld, N. (1995). Does it hurt to say um? *Journal of Nonverbal Behavior, 19*, 171–186.

Corriveau, K. H., Meints, K., & Harris, P. L. (2008). Early tracking of informant accuracy and inaccuracy. *British Journal of Developmental Psychology, 27*, 331–342.

Derwing, T. M., & Munro, M. J. (2009). Putting accent in its place: Rethinking obstacles to communication. *Language Teaching, 42*, 476–490. doi:10.1017/S02614448 0800551X

Fox Tree, J. E., & Clark, H. H. (1997). Pronouncing 'the' as 'thee' to signal problems in speaker. *Cognition, 62*, 151–167.

Harris P.L., and Corriveau, K.H., (2011). Young children's selective trust in informants. *Philosophical Transactions of the Royal Society, 366,* 1179-1187.

Hirschfeld, L., & Gelman, S. (1997). What young children think about the relationship between language variation and social difference. *Cognitive Development, 12*, 213–238.

Horst, J. S. & Hout, M. C. (2014). The Novel Object and Unusual Name (NOUN) Database: a collection of novel images for use in experimental research. *Unpublished manuscript.*

Jaswal, V. K., & Malone, L. S. (2007). Turning believers into skeptics: 3-year-olds' sensitivity to

cues to speaker credibility. *Journal of Cognition and Development, 8*(3), 263-283.

Jaswal, V. K., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science, 17(9*), 757–758.

Kidd, C., White., K., Aslin., R (2011). Toddlers use speech disfluencies to predict speakers' referential intentions. *Developmental Science 14*(4), *925-934*

Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children's selective trust in native-accented speakers. *Developmental Science, 14*(1), 106-111.

Kinzler, K.D., Dupoux, E., & Spelke, E.S. (2007). The native language of social cognition. *Proceedings of the National Academy of Sciences of the United States of America, 104*, 12577–12580.

Koenig, M. A., & Harris, P. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development, 76*, 1261–1277.

Lane, J. D., and Harris, P. L., (2015). The Roles of intuition and informants' expertise in children's epistemic trust. *Child Development, 86*(3): 919-926.

Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development, 73*, 1073–1084.

MacDonald, K., Schug, M., Chase, E., & Barth, H. (2013). My people, right or wrong? Minimal group membership disrupts preschoolers' selective trust. *Cognitive Development, 28*(3), 247–259.

Moore, C., Bryant, D., & Furrow, D. (1989). Mental terms and the development of certainty. *Child Development, 60*, 167–171. doi: 10.2307/1131082

Orena, A., & White., K., (2015). I Forget what that's called! Children's online processing of disfluencies depends on speaker knowledge. *Child Development, 86*(6), 1701-1709.

Owens, S., and Graham, S., (2016). Thee, uhh disfluency effect in preschoolers: A cue to discourse status. *British Journal of Developmental Psychology, 34,* 388-401.

Rakoczy, H., Hamann, K., Warneken, F., and Tomasello, M. (2010). Bigger knows better: Children selectively learn rule games from adults rather than from peers. *British Journal of Developmental Psychology 28,* 785-798.

Sabbagh, M. A. & Baldwin, D.A., (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development, 71(*4), 1054-1070.

Sabbagh, M.A., & Shafman, D., (2009). How children block learning from ignorant speakers. *Cognition 112,* 415-422.

Schacter, S., Christenfeld, N., Ravina, B., & Bilous, F. (1991). Speech disfluency and the structure of knowledge. *Journal of Personality and Social Psychology, 60,* 362–367.

Scott, K., Chu, J., and Schulz, L., (2017). Lookit (Part 2): Assessing the viability of online developmental research, Results from three case studies. *Open Mind: Discoveries in Cognitive Science, 1*(1), 15–29. doi:10.1162/opmi_a_00001

Sobel, D. M., & Finiasz, Z. (2020). How children learn from others: An analysis of selective word learning. *Child Development, 91*(6), e1134-e1161.

Taylor M, Cartwright BS, 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), 1334-1351.

Tenney, E. R. A., Small, J. E., Kondrad, R. L., Jaswal, V. K., & Spellman, B. (2011). Accuracy, confidence, and calibration: How young children and adults assess credibility. *Developmental Psychology, 47,* 1065–1077.

Von Tiling, J. (2011). Listener perceptions of stuttering, prolonged speech, and verbal avoidance behaviors. *Journal of Communication Disorders*, *44*, 161–172.

VanderBorght, M., and Jaswal, V.K., (2009). Who knows best? Preschoolers sometimes prefer child informants over adult informants. *Infant and Child Development 18*(1), 61-71.

Wimmer H, Hogrefe GJ, Perner J., (1988). Children's understanding of informational access as source knowledge. *Child Development, 59*, 386–396.

Weatherhead, White, & Friedman (2016). Where are you from? Preschoolers infer background from accent. *Journal of Experimental Child Psychology 143*, 171–178.

Wheatherhead, White & Friedman (2018). Accent, Language, and Race: 4–6-year-old children's inferences differ by speaker cue. *Child Development 89(5),* 1613–1624.

White, W., Nilsen, E., Deglint, T., Silva, J., (2019). That's thee, uhh blicket! How does disfluency affect children's word learning? *First Language, 40*(1), 3-20.

Zhang, M., and Sylva, K., (2021). Effects of group membership and visual access on children's selective trust in competitive and non-competitive contexts. *Cognitive Development, 57.*

Appendix A

E1 Procedure

Audio Check

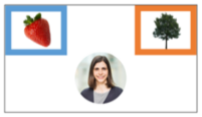

Trial 1

Familiarization Phase

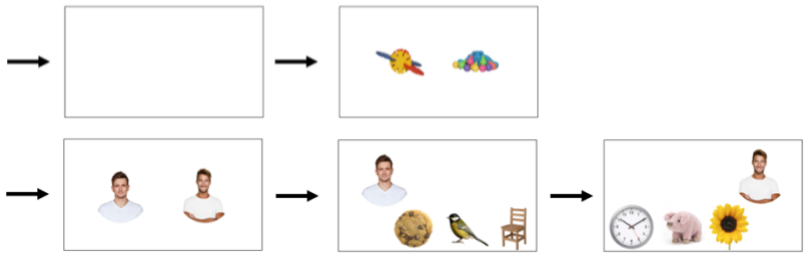

Test Phase

Trial 2

Familiarization Phase

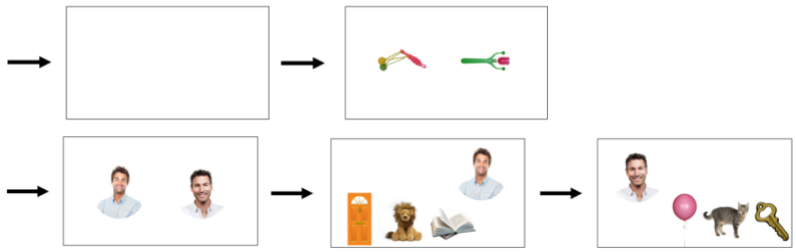

Test Phase

Explicit Judgement Question

E2 Procedure

Audio Check

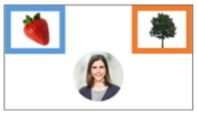

Trial 1

Familiarization Phase

Test Phase

Trial 2

Familiarization Phase

Test Phase

Explicit Judgement Question

Appendix B

Child Accent Ratings: Procedure

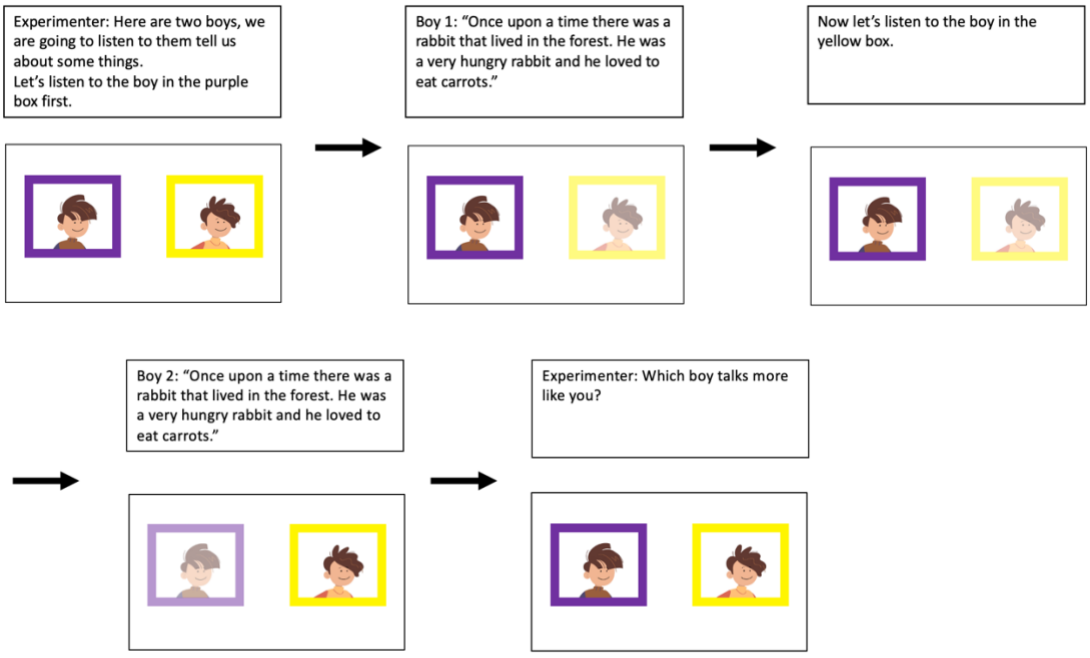

Appendix C

E1 Script

Audio Check

Experimenter: First, we are going to make sure you can hear everything okay. Let's listen to the girl!

Girl: Can you tell me what colour box the strawberry is in?
Girl: Can you tell me what colour box the tree is in?

Trial 1

Familiarization

Experimenter: Awesome job! Okay so today we're going to play a game where we have to figure out the names of some new toys!

Experimenter: [two novel objects appear] Here are two new toys. We have to figure out what one of these toys is called.

Experimenter: [still images of two male character appear] Here are two guys. We're going to listen to them tell us about some thinfs, and then you can pick one of them to help us figure out the name of the new toys!

Speaker 1 [flashes]: Hello, let's play a game.
Speaker 2 [flashes]: Hi there, yes, let's play.

Experimenter: First, let's listen to this guy!

Speaker 1: That is a cookie. That is a bird. That is a chair. [objects flash]

Experimenter: Now let's listen to this guy!

Speaker 2: That is a uhmm clock. That is a uhmm pig. That is a uhmm flower. [objects flash]

Test Phase

Experimenter: Remember these toys? One of these is a WILP! Let's listen to what the guys say!

Speaker 1: This is a wilp [toy bellow speaker 1 flashes]

Speaker 2: This is a uhmm wilp [toy bellow speaker 2 flashes]

Experimenter: Which one is a WILP? The one in the blue box or the one in the orange box?

Trial Two

Familiarization Phase

Experimenter: Now we are going to play again!

Experimenter: [two new novel objects appear] Here are two new toys. We have to figure out what one of these toys is called.

Experimenter: [still images of two new male character appear] Here are two more guys. We're going to listen to them tell us about some things, and then you can pick one of them to help us figure out the name of the new toys!

Speaker 2 [flashes]: Hello, let's play a game.
Speaker 1 [flashes]: Hi there, yes, let's play.

Experimenter: First, let's listen to this guy!

Speaker 2: That is a uhmm door. That is a uhmm lion. That is a uhmm book. [objects flash]
Experimenter: Now let's listen to this guy!

Speaker 1: That is a balloon. That is a cat. That is a key. [objects flash]

Test Phase

Experimenter: Remember these toys? One of these is a mense! Let's listen to what the guys say!

Speaker 1: This is a mense [toy bellow speaker 1 flashes]
Speaker 2: This is a uhmm mense [toy bellow speaker 2 flashes]

Experimenter: Which one is a mense? The one in the blue box or the one in the orange box?

Explicit Judgement Question

Experimenter: Okay, we are almost done! I just have one last question for you. Was there anything funny about the way the guys were talking? What? [If *child notices disfluency, ask why they think the speaker was disfluent → e.g., "why do you think they were saying 'umm'?"]*

E2 Script

Experimenter: First, we are going to make sure you can hear everything okay. Let's listen to the girl!

Girl: Can you tell me what colour box the strawberry is in?
Girl: Can you tell me what colour box the tree is in?

Trial 1

Familiarisation Phase

Experimenter: Awesome job! Okay so today we're going to play a game where we have to figure out the names of some new toys!

Experimenter: [two novel objects appear] Here are two new toys. We have to figure out what one of these toys is called.

Experimenter: [still images of two male character appear] Here are two guys. We're going to listen to them tell us about some things, and then you can pick one of them to help us figure out the name of the new toys! You can only ask one of the guys to help us, so make sure you listen carefully!

Experimenter: First let's listen to the guy in the blue box.

Speaker 2 [flashes]: Hello, let's play a game.
Experimenter: Now let's listen to the guy in the orange box!
Speaker 1 [flashes]: Hi there, yes, let's play.

Experimenter: First, let's listen to the guy in the blue box!

Speaker 1: That is a cookie. That is a bird. That is a chair. [objects flash]

Experimenter: Now let's listen to the guy in the orange box!

Speaker 2: That is a uhh clock. That is a uhh pig. That is a uhh flower. [objects flash]

Test Phase

Experimenter: Remember these toys? One of these is a WILP! Let's listen to what the guys say!

Speaker 1: This is a wilp [toy bellow speaker 1 flashes]
Speaker 2: This is a uhh wilp [toy bellow speaker 2 flashes]

Experimenter: Which one is a WILP? [If child points: What colour box is that one in?]

Trial Two

Familiarization Phase

Experimenter: Now we are going to play again!

Experimenter: [two new novel objects appear] Here are two new toys. Now, we have to figure out what one of these toys is called.

Experimenter: [still images of two new male character appear] Here are two more guys. We're going to listen to them tell us about some things, and then you can pick one of them to help us figure out the name of the new toys! Remember, you can only ask one of the guys to help us, so make sure you listen carefully!

Experimenter: First, let's listen to the guy in the blue box.

Speaker 2 [flashes]: Hello, let's play a game.

Experimenter: Now let's listen to the guy in the orange box!

Speaker 1 [flashes]: Hi there, yes, let's play.

Experimenter: First, let's listen to the guy in the blue box!

Speaker 2: That is a uhh door. That is a uhh lion. That is a uhh book. [objects flash]
Experimenter: Now let's listen to the guy in the orange box!

Speaker 1: That is a balloon. That is a cat. That is a key. [objects flash]

Test Phase

Experimenter: Remember these toys? One of these is a mense! Let's listen to what the guys say!

Speaker 1: This is a mense [toy bellow speaker 1 flashes]
Speaker 2: This is a uhh mense [toy bellow speaker 2 flashes]

Experimenter: Which one is a mense? [If child points: What colour box is that one in?]

Explicit Judgement Question

Experimenter: Okay, we are almost done! I just have one last question for you. Was there anything funny about the way the guys were talking? What? [If *child notices disfluency, ask why they think the speaker was disfluent → e.g., "why do you think they were saying 'umm'?"]*

Appendix D

Familiar Objects









Novel Objects