Preschoolers' Understanding of Graded Preferences

Jane Hu¹, Christopher G. Lucas², Thomas L. Griffiths³, Fei Xu³

¹ University of Washington, USA; janehu@uw.edu

² University of Edinburgh, UK; c.lucas@ed.ac.uk

³ University of California, Berkeley, USA; [\[fei_xu\]](mailto:[fei_xu]), [\[tom_griffiths\]@berkeley.edu](mailto:[tom_griffiths]@berkeley.edu)

Corresponding author:

Jane Hu
University of Washington
Institute of Learning and Brain Sciences
Portage Bay Building
Box 357988
Seattle, WA 98195, USA
+001 (206) 221-1684
janehu@uw.edu

UNDERSTANDING GRADED PREFERENCES

Abstract

To navigate the social world, children must learn about others' preferences. Though people can use emotional and verbal cues to express their preferences, these cues are often unavailable or unreliable. Previous research has found that preschoolers and toddlers use statistical information to infer the existence of a preference. However, in the real world, preferences are not binary; they can also be graded. In two experiments, we find that preschoolers use statistical information about an agent's choices to infer the graded strengths of preferences. From observing an agent's choices, preschoolers inferred that objects the agent chose less consistently were less preferred than objects the agent chose more consistently. Additionally, preschoolers' responses indicated that preschoolers make more sophisticated transitive inferences than previously attributed to this age group.

Keywords: preferences; social cognition; statistical inference; transitive inference; cognitive development.

Preschoolers' Understanding of Graded Preferences

Inferring others' mental states from their actions is an indispensable social skill. As adults, we often observe others' actions to infer what they like or dislike. For example, if our colleague Jim buys an orange soda every day from a vending machine with many other options available, we might infer that Jim has a preference for orange soda. Because preferences are relatively stable dispositions, identifying a person's preferences can help us predict how an agent may behave in the future (e.g., Jim will likely buy orange soda again), how to become a better social partner (e.g., buying Jim an orange soda as a thank-you for a favor), or how much we'll like something new based on our shared preferences with that person (Fawcett & Markson, 2010).

How does the ability to infer preferences develop? One prerequisite for inferring preferences is the recognition that other people are intentional agents. Infants as young as 3 months old can interpret agents' reaching behavior as goal-directed (Woodward, 1998; Sommerville, Woodward, & Needham, 2005), and 6-month-olds take into account others' perceptual access when making judgments about agents' preference for objects they consistently reach towards (Luo & Baillargeon, 2007; Luo & Johnson, 2009). By 12 months, infants track multiple agents' behaviors and understand that goals are person-specific (Buresh & Woodward, 2007).

To infer others' preferences, children must also recognize what cues are informative, and when. Previous studies suggest toddlers and preschoolers infer preferences from emotional or verbal cues (Lumeng, Cardinal, Jankowski, Kaciroti, & Gelman, 2008; Repacholi & Gopnik, 1997). However, these cues are not always available; people do not always wear their emotions on their sleeves. Even worse, these cues could be unreliable. Our emotional states aren't always a

result of our choices, and could be misleading if taken to represent our attitudes towards our choices. For instance, a coworker who appears irritated as she buys coffee is probably not irritated by her choice of beverage – instead, she may have just had a bad morning.

In the case of inferring preferences, the *choices* an agent makes are often more reliable cues than their emotional or verbal responses. Actions often reflect preferences, as people tend to choose options they like or avoid options they don't. Recent studies suggest children use statistical evidence to infer preferences; preschoolers and 20-month-old toddlers infer that a puppet has a preference when its choices are inconsistent with random sampling (Kushnir, Xu & Wellman, 2010; Ma & Xu, 2011). Preschoolers also use non-random sampling and multiple agents as cues to the generalizability of preferences (Diesendruck, Salzer, Kushnir & Xu, 2014).

Children's ability to identify a preference for one object over another – i.e. an agent likes object X more than object Y, or Y more than X – is well established. However, inferring the gradedness of preferences is more complex. When inferring the existence of a preference, one can succeed by using a binary heuristic – for instance, violation of randomness – as cues, but inferring graded preferences requires an extra step: comparing the *degree* of preference of multiple options to one another.

Previous developmental research suggests children may have some key skills to help them recognize and compare graded preferences. Five-year-olds have been found to understand scalar implicature – specifically, they are able to recognize the distinction between “all” and “some” (Papafragou & Musolino, 2003; Papafragou & Tantalou, 2004). Research on preschoolers' transitive inference abilities also suggests that children may have some of the elementary skills necessary for inferring graded preferences. The earliest transitive inference studies found that 4-year-old children struggle with transitive inference in word problems (Piaget, 1928; 1955), but

this failure could have been due to memory limitations. Other studies found that 4-year-olds can make transitive inferences about spatial position with help from visual cues like ordered rods or block towers to aid memory (Bryant & Trabasso, 1971; Halford, 1984; Pears & Bryant, 1990; Andrew & Halford, 1998).

More recently, work by Mou, Province, & Luo (2014) found that infants demonstrate transitive reasoning abilities about others' preferences: if an agent repeatedly reaches for A over B and B over C, then 16-month-olds show surprise (i.e., longer looking time) if the agent reaches for C instead of A. Even children as young as 9 months show surprise when an experimenter grasps an object inconsistent with the hierarchy of choices she previously demonstrated (Robson et al., 2014). Other research suggests that these abilities could extend to the social domain. Mascaro & Csibra (2014) found 15-month-olds make incremental inferences about dominance relationships. While their study does not directly address the question of transitive inference, making incremental inferences is an essential skill for recognizing broader social structures. Interestingly, their data found children had an easier time recognizing linear relationships than circular ones, another skill that would assist with children's transitive inference abilities.

Children's performance in causal learning tasks also suggests they can make transitive inferences about causes and effects. Children as young as three years have demonstrated an understanding that if X causes Y and Y causes Z, then X causes Z, and five-year-olds can explicitly state the necessity of Y in the relationship between X and Z (Shultz, Pardo, & Altmann, 1982). Furthermore, Schulz, Gopnik and Glymour (2007) found that preschoolers successfully identify transitive relationships between causes. Children were introduced to an electronic toy that contained a switch and two gears, and watched an experimenter intervene on the toy (i.e.

turning the switch on or off, or removing one of the gears) to infer the causal relationships between the components of the toy. Depending on the result of the experimenter's interventions, children inferred different transitive causal relationships between the components – for instance, that the switch causes Gear 1 to spin, which causes Gear 2 to spin, versus the Switch causes Gear 2 to spin, which causes Gear 1 to spin.

However, transitive inference in causal relationships differs from transitive inference in preferences in a crucial way: the relationships between causes change with the omission of a causal element, whereas the relationships between preferences do not. To use the Schulz, Gopnik, and Glymour study as an example, if the switch causes Gear 1 to spin, which causes Gear 2 to spin, Gear 1 is necessary for the transitive relationship between the Switch and Gear 2; i.e., without Gear 1, turning on the Switch may not cause Gear 2 to spin. In transitive inferences about preferences, however, the existence of an intermediate preference doesn't change the relationship between other preferences. For instance, if one prefers X most, Y next most, and Z least, one could still infer that X is preferred to Z without any knowledge of where Y falls on the spectrum of preferences.

The task we used in the current studies asked children to infer an agent's graded preferences from their choices, and further investigated children's transitive inference abilities in a social domain. Experiment 1 examined the inferences children make about an agent's preferences after observing his choices. This task required children to integrate information from two relational premises – the agent's choices between objects A and C, and B and C – to make inferences about an agent's overall preferences for A, B, and C by asking children to predict which of those objects would be preferred to a novel object D. In demonstrations, children saw that A was very consistently chosen over C, whereas B was only somewhat consistently chosen

over C. Unlike other transitive inference tasks which show $A > B$ and $B > C$, children could not use adjacent pairings to infer $A > B > C$; rather, they had to make inferences about the relative preferences of A and B *relative* to C.

Moreover, the experiment was designed to elucidate what strategies children were using to evaluate the agent's preferences: the agent was faced with the A/C pairing and B/C pairing different numbers of times, so that children would arrive at different inferences about the agent's preference. An easy heuristic to use would be for children to infer the agent had a preference for the object chosen the most absolute number of times; the more sophisticated strategy would be to infer a preference for objects based on the proportion of times the object was chosen.

1. Experiment 1: Inferring graded preferences

In Experiment 1, participants watched as Duckie, a puppet, chose one of two objects. When object A and object C were presented, Duckie chose A 5 out of 5 times (100%). When objects B and C were presented, Duckie chose B 7 out of 10 times (70%). Thus overall, object A was chosen 100% of the time over competitors, object B was chosen 70% of the time, and object C was only chosen 20% of the time (3 out of 15 times), consistent with a hierarchy of preferences where A is the most preferred, B is the second most preferred, and C is the least preferred (see Figure 1 for a schematic representation of the procedure). Duckie's choices were designed to control for mere association effects; in order to succeed at inferring this hierarchy of preferences, children must infer that the relative proportion of Duckie's choices, not absolute number of choices, indicate his preferences, and there is no direct evidence that $A > B > C$; rather, children must make inferences about the relative strength of preferences for A vs. B by comparing Duckie's choice rates of A over C and B over C.

Children were asked to make inferences about Duckie's preferences in two test questions. In the first test question, participants were asked to give Duckie the object they believed he would like the most: object A (chosen in 100% of the trials it appeared in) or B (chosen in 70% of the trials it appeared in). If children tracked the proportion of Duckie's choices for A and B in the demonstration phase, they should infer that Duckie prefers object A over B. In contrast, if they inferred preferences by the absolute number of times the objects were chosen, they would infer Duckie preferred B over A.

In the second test question, participants were asked to judge Duckie's preferences for objects A, B, and C compared to a novel object D. Participants were asked to give Duckie the object he preferred from the following pairs: A / D; B / D; and C / D. Given that Duckie did not make any choices that involved D, children should view D as a neutral object; if children inferred that Duckie's choices of A, B, and C represent graded preferences, this should influence the rates of their choices for A, B, and C compared to D. Specifically, they should infer that strongly preferred objects would be preferred to D, whereas dispreferred objects would continue to be dispreferred relative to D. Therefore, children should believe that Duckie is very likely to prefer object A over D, somewhat likely to prefer B over D, and unlikely to prefer C over D.

1.1 Method

1.1.1 Participants. Participants were 31 preschoolers (mean age = 4 years 5 months; range = 44 – 63 months; 14 female). Participants were recruited in a major metropolitan city by mail and phone calls or from local preschools, and were predominantly Caucasian and middle class. An additional two children were tested, but were uncooperative and thus excluded.

1.1.2 Materials and Procedure. A set of four novel objects was used in the study (see Figure 1); which objects were designated as A, B, C, or D were counterbalanced.

To minimize novelty effects, the study began with a familiarization phase where children played with the four objects. If children did not spontaneously play with them, the experimenter picked up each object, presented it to the child, and asked, “Did you see this one?”

A strong positive or negative preference for any specific object could affect children’s test responses, so each child was asked a baseline question to assess their preferences. The experimenter presented each participant with two objects randomly selected from the full set of four stimuli objects. The experimenter then asked, “Which one do *you* like more?” Whether the baseline question was presented before the demonstration phase or after test questions was counterbalanced.

Next, children watched as Duckie made his choices in two blocks of demonstrations: a 70% versus 30% block and a 100% versus 0% block. The order of trial block presentation was counterbalanced. In the 70% versus 30% demonstration block, each child sat at a table across from the experimenter. The experimenter asked the child, “Should we invite my friend Duckie to play with us?” and introduced the puppet. Then, the experimenter placed objects B and C on the table, approximately 18 inches apart, and said, “Let’s ask Duckie which one he wants to play with.” Duckie appeared from beneath the table and stood behind it, equidistant from the objects. Before each of Duckie’s choices, the experimenter asked, “Duckie, which one?” Duckie looked back and forth between the objects as if examining his options, then picked one up, held it for one second, and placed it back on the table before returning to his spot behind the table. This was repeated for a total of ten trials. In seven of the ten trials, Duckie chose object B. In the other three trials, Duckie chose object C. The side on which each object appeared was counterbalanced across participants, and the order of Duckie’s looks between objects B and C and the order in which he made his choices were randomized.

The 100% versus 0% demonstration block was the same as the 70% versus 30% block, except that Duckie chose between objects A and C, and made only five choices. Duckie chose object A all five times.

Finally, children were asked two test questions. In Test Question 1, participants were presented with objects A and B and asked, “Which one does Duckie like more?”

In Test Questions 2a, 2b, and 2c, participants were asked about Duckie’s preferences for each of the objects in the demonstration phase (A, B, and C) compared to a novel object D. For each comparison, participants were asked, “Can you give Duckie the one he likes more?” The order in which 2a, 2b, and 2c were asked was randomized.

1.2 Results

Binomial tests revealed that no objects were selected at higher-than-chance rates on the baseline trials, suggesting that there were no sample-wide patterns in participants’ preferences for stimulus objects. Table 2 presents children’s baseline responses in both experiments. Fisher’s exact tests found no effects of sex, demonstration block (70% v. 30% block or 100% v. 0% block first), or the order of questions 2a, 2b, and 2c in responses to test questions.

Figure 2 shows the main results of Experiment 1. In Test Question 1, 25 of 31 children (81%) chose object A over B, binomial test, $p < .001$. In Test Question 2, 24 of 31 children (77%) chose object A over the novel object D, 19 of 31 children chose (61%) object B over the novel object D, and 12 of 31 children (39%) chose object C over the novel object D. A Cochran’s Q test revealed a statistically significant difference in the rate of selecting the familiar objects (A, B, C) over the novel object D in questions 2a, 2b, and 2c, $\chi^2(2, N = 31) = 8.07, p < .01$.

We also examined children's patterns of responses for additional evidence that *individual* children inferred the gradedness of Duckie's preferences. Since children were asked for responses to binary questions, data from individual children can reveal indirect yet useful clues about children's inferences. If children did indeed infer that $A > B > C$, they would have chosen the 100% object over the 70% object in Test Question 1 (25 of 31 children). Additionally, their choices in Test Questions 2a, 2b, and 2c should reflect this belief; they should infer that A is at least as likely as B to be preferred over the novel object, and that B should be at least as likely as C to be preferred over the novel object. There are a total of eight possible patterns children's responses in Test Question 2 could take (see Table 2 for a full list), and there are four patterns that could reflect a belief that $A > B > C$. One would be choosing the 100% object in 2a, the novel object in 2b, and novel object in 2c; if Duckie prefers A the most, he should choose it over a novel object, but would still prefer the novel object over B or C. Another pattern would be to choose the 100% object in 2a, the 70% object in 2b, and the novel object in 2c; this would reflect an understanding that C is least preferred, given that the novel object was chosen over it but not A or B. A pattern of choosing all familiar objects (100%, 70%, 20%) or all novel objects could also reflect this belief.

Of the 25 children who chose the 100% object over the 70% object in question 1, 18 responded with a pattern in question 2 that suggests they inferred $A > B > C$ (binomial test $p < .05$; see Table 2). In an analysis including all participants (including those who chose the 70% object over the 100% object in question 1), 20 of 31 children showed one of the patterns ($p = .14$; see Table 3).

1.3 Discussion

Children's choices on test questions suggest they used Duckie's choices to infer a hierarchy of graded preferences with object A (chosen 100% of the time over competitors) being the most preferred, followed by B (chosen 70% of the time over competitors) and C (chosen 20% of the time over competitors). As a group, children's responses indicate that they inferred that Duckie preferred A over B. Furthermore, children inferred that the likelihood Duckie would prefer a novel object over an object shown in demonstration was graded with respect to the consistency with which Duckie chose it during demonstrations. Because Duckie chose A consistently, children inferred he would likely continue to prefer it over a novel object, whereas he was only somewhat likely to prefer B over the novel object, and unlikely to prefer C over the novel object.

These results also suggest that children can make transitive inferences based on their observations. In order to succeed in this experiment, children had to integrate several relational premises to make inferences about an agent's graded preferences. To form a mental hierarchy that reflects these graded preferences, children had to first observe each set of choices to make inferences about their relative value (e.g., A is strongly preferred over C, and B is somewhat preferred over C), and then integrate those relational premises to make inferences about the choices' relative values: $A \gg C$ and $B > C$, therefore $A > B$.

A possible alternative explanation for these results is that children used a heuristic to answer test questions without understanding that Duckie's preferences were *graded*. Rather than tracking the consistency of Duckie's choices, children could have used a simpler heuristic to infer Duckie's preferences. For instance, children could have inferred that object A was strongly preferred because it was chosen *all* of the time, whereas objects B and C were less preferred because they were chosen only *some* of the time. Children's responses to Test Questions 2b

(12/31 children chose novel object) and 2c (19/31 chose novel object) suggest that they inferred the gradedness of preferences when comparing objects chosen *some* of the time, but this issue could be more thoroughly explored. Experiment 2 was designed to address this issue, and to extend and replicate the results of Experiment 1.

2. Experiment 2: Inferring graded preferences for objects chosen some of the time

In Experiment 2, we tested preschoolers with a procedure similar to the one used in Experiment 1, except that during demonstrations, Duckie chose all objects only *some* of the time. If children rely on a simple heuristic like inferring objects chosen all of the time are preferred over objects chosen only some of the time, they would fail at this task. Alternatively, children's success would provide further support that they are indeed tracking the proportion, not absolute number, of Duckie's choices to infer his graded preferences.

2.1 Method

2.1.1 Participants. Participants were 35 preschoolers (mean age = 4 years 1 month; range = 42 – 62 months; 24 girls). Participants were recruited from a major metropolitan city by mail and phone calls or from local preschools, and were predominantly Caucasian and middle class. An additional eight children were tested, but excluded due to experimenter error (6), video malfunction (1), and uncooperativeness (1).

2.1.2 Materials and Procedure. The same materials and general procedure from Experiment 1 were used in Experiment 2. The only difference between the two procedures was the set of choices Duckie made during the two demonstration blocks.

The 63% versus 37% demonstration block was identical to the first demonstration block in Experiment 1, except Duckie made a total of 11 choices. In 7 of the 11 trials (63%), Duckie chose object B. In the other four trials, Duckie chose object C.

The 83% versus 17% demonstration block was identical to the second demonstration block in Experiment 1, except that Duckie made a total of 6 choices. Duckie chose object A 5 out of 6 times (83%), and C 1 out of 6 times (17%).

2.2 Results and Discussion

Binomial tests revealed that no objects were selected at higher-than-chance rates on the baseline trials, suggesting that there were no sample-wide patterns in participants' references for stimulus objects. Fisher's exact tests found no effects of sex, demonstration block (63% v. 37% or 83% v. 17% block first), or the order of questions 2a, 2b, and 2c in responses to test questions.

Figure 3 shows the main results of Experiment 2. In Test Question 1, 24 of 35 children (69%) chose object A over B, binomial test, $p = .04$. In Test Question 2, 25 of 35 children (71%) chose object A over the novel object D, 18 of 35 children chose (51%) object B over the novel object D, and 13 of 35 children (37%) chose object C over the novel object D. A Cochran's Q test revealed a statistically significant difference in the rate of selecting the familiar objects (A, B, C) over the novel object in questions 2a, 2b, and 2c, $\chi^2(2, N = 35) = 8.34, p < .01$.

Like in Experiment 1, children's patterns of choices in Experiment 2 suggest that individual children inferred the gradedness of Duckie's preferences. If children inferred that Duckie's preferences were $A > B > C$, they would have chosen the 83% object over the 63% object in question 1, and in question 2, they should infer that A is more or just as likely to be preferred over the novel object than B, *and* that B should be more or just as likely to be preferred over the novel object than C. The patterns that would reflect this belief are: 83% over novel object (2a), 63% over novel object (2b), novel object over 29% object (2c); 83%, novel, novel; 83%, 63%, 29%; novel, novel, novel.

Of the 24 children who chose the 83% object over the 63% object in question 1, 20 children's responses fit one of those patterns (binomial test $p < .001$; see Table 4). In an analysis including children who chose the 63% object over the 83% object in question 1, 26 of 35 children's responses matched of the patterns (binomial test $p < .01$; see Table 5).

3. General Discussion

In two experiments, we found that children observed an agent's choices to infer his graded preferences. Children inferred that Duckie preferred an object chosen 100% of the time over an object chosen 70% of the time, or an object chosen 83% of the time over an object chosen 63% of the time. Children also used Duckie's choices to make inferences about the strength of his preferences; their responses suggest that they inferred that objects chosen less consistently were less preferred and less likely to be preferred compared to a novel object, whereas objects chosen consistently were more preferred and more likely to be preferred compared to a novel object.

These studies demonstrate that young children use statistical evidence to infer graded preferences using indirect transitive inference. Without explicit emotional or verbal cues, children can use statistical evidence from choice actions to infer the agent's preferences, and use those inferences to make predictions about what an agent would prefer in the future. Children do not simply represent preferences as all or none; instead they can use statistical evidence to establish a hierarchy of preferences.

These results also have important implications for the study of children's ability to make transitive inferences. Whereas previous work suggests that 4-year-olds would succeed in basic transitive inferences ($A > B$, $B > C$, $\therefore A > C$), children's success in these experiments indicate competence in making more sophisticated, indirect transitive inferences. In our studies, objects A

and B were never directly compared. The only available information about the relationship between A and B was each of their relationships with object C: Duckie chose A more consistently over C ($A \gg C$), whereas Duckie chose B only somewhat consistently over C ($B > C$).

These studies also suggest new questions for future investigation. First, the mechanism children are using to infer Duckie's preferences could be further explored. There are several strategies children could use to infer the relative strengths of preferences. An easy strategy would be to track the *raw frequency* of an agent's choices, and to assume that the number of times an agent chooses an option has a direct relationship to the agent's preference strength. For instance, if A is chosen five times and B is chosen seven times, then the agent has a stronger preference for B than for A. Previous work has found that infants succeed at tracking frequency information in linguistic input, even in noisy environments (Smith & Yu, 2007; Vouloumanos & Werker, 2009), and preschoolers track the frequency of a person's positive or negative behaviors to predict their future behaviors and infer their personality traits (Boseovski & Kang, 2006). These studies suggest preschoolers should be adept at tracking the frequencies of different choices.

However, this set of studies suggests that children are not only tracking frequency of choices. Even though children were shown that Duckie chose object B a greater *number* of times than A (seven times versus five in Experiment 1, or seven times versus four in Experiment 2), children still inferred that Duckie preferred A over B.

A more sophisticated strategy for determining preference strength would require children to go beyond raw frequency counts and consider the agent's choices relative to the total number of choices made. The number of times an option is chosen only matters in the greater context of how many total choices there were; the *proportion* of times an object is chosen could indicate preference strength. Indeed, in this study, children's responses indicate that they inferred that

Duckie preferred toys he chose a high proportion of times over toys that he chose a lower proportion of times; for example, they inferred he preferred object A (chosen 100% or 83% over competitors) over object B (chosen 70% or 63% over competitors).

However, one possible alternate explanation is that rather than tracking the proportion of Duckie's choices, children could instead track the frequency with which Duckie *rejects* a toy. For instance, in Experiment 2, children could have inferred that object C was least preferred because it was rejected the most often (12 out of 17 choices), that object B was next least preferred (rejected 4 out of 11 choices), and that object A was most preferred because it was rejected only once out of 6 choices. To our knowledge, there have been no published studies that suggest children track the frequency of rejected items in choice tasks, but the procedure of these two experiments does not allow us to rule out the possibility that children could be using frequency of rejections as a cue to Duckie's preferences. Further investigation is necessary to determine whether children in this study were tracking Duckie's proportion of choices or frequency of rejections, as both strategies could lead children to infer an $A > B > C$ hierarchy in Experiments 1 and 2.

Yet another alternate explanation could be that children were not reasoning about Duckie's preferences, but rather, they assigned values to objects based on its associations with positivity or negativity. Researchers studying pigeons' transitive inference abilities introduced the value transfer theory (VTT): that animals' apparent inference skills can be explained by stimuli's values, which are based on reinforcement and/or association with other stimuli (von Fersen et al, 1991; Zentall & Sherburne, 1994). In our study, it is possible that in Test Question 1, children chose object A over B not because they tracked Duckie's preferences, but because A was associated with the rejected object C fewer times than B was. (A was paired with C 5 times

in Experiment 1 and 6 times in Experiment 2, whereas B was paired with C 10 times in Experiment 1 and 11 times in Experiment 2.) However, the VTT cannot account for data from Test Question 2, in which children were asked to compare each object seen in the demonstration phase – objects A, B, and C – with a novel object D. VTT would predict that D – which has no value assigned, since it has not been compared to anything – should be preferable to A, which has been associated with the “loser” object C several times.

Additionally, in these studies, Duckie’s choices indicated his preferences, but other factors, like the amount of data observed, can also affect inferences about preference. We might infer that choosing an object 99 out of 100 times (99%) reflects a stronger preference than choosing it 1 out of 1 time (100%), even though it was chosen more consistently in the latter case. The number of observations affects our confidence in making predictions about preferences. In the 1 out of 1 case, we have very little information with which to predict the strength of the agent’s preference and whether (s)he will continue to prefer this object in future choices. However, in the 99 out of 100 case, we have much more information, and the agent’s choices give us more confidence that (s)he strongly prefers the object chosen. Recent Bayesian modeling work suggests that the amount of available information should influence inferences about preferences (Lucas et al., 2014), and future work can investigate whether children share this intuition.

It could also be worth investigating whether our assumptions about object D’s neutrality were an accurate description of children’s views. Although we believe our results are consistent with the hypothesis that children viewed D as a neutral object, children nonetheless could have made different assumptions about D that were not discernable by the current design. For instance, children may not view object D as neutral, but may ascribe a slightly positive value to it

by virtue of its novelty, or a slightly negative value to it due to the lack of information about its value to Duckie. Results from such an investigation could have implications for other studies that use a similar methodology.

While the current studies focused on children's use of statistical evidence to make inferences about positive preferences, there has been little investigation into how children infer negative preferences (i.e., dislike). Children interpret choice actions as evidence for a preference. Would they interpret lack of choice as evidence for a negative preference? In our studies, it is unclear whether children interpreted Duckie's relatively few choices of object C as a very weak preference, or a dispreference. While there is evidence that adults can use choice behavior to infer the strengths of negative preferences (Jern, Lucas, & Kemp, 2011), it is not yet clear when this ability develops.

Furthermore, even if choices are indicative of preferences, it may be difficult to evaluate and compare these preferences if they involve choices from different categories. In these studies, Duckie was faced with options from the same category – all options were toys. Had the options been pulled from different categories – say, toys, snacks, and tools – it would be more difficult for participants to evaluate the constancy of preferences across these categories. Future studies could elucidate the strategies children use when faced with making these types of inferences.

Lastly, choices are not always indicative of a positive or negative preference. Positive and negative preferences involve positive or negative appraisals, whereas choices can be indicative of a neutral internal mental state, such as a goal. For example, imagine that you observe a colleague picking up a stapler. You are unlikely to infer that she has a preference for staplers. Rather, you might infer that she has a goal of stapling paper together. Future studies

may examine when children develop the understanding that choices may reflect different mental states such as goals, preferences, and beliefs.

Acknowledgments

This research was supported by a National Science Foundation Graduate Research Fellowship to JH. We would also like to thank all participants and their parents for supporting our research, and the members of the Berkeley Early Learning Lab for their helpful suggestions.

References

- Andrews, G., & Halford, G. S. (1998). Children's ability to make transitive inferences: The importance of premise integration and structural complexity. *Cognitive Development*, 13(4), 479-513.
- Bryant, P. E., & Trabasso, T. (1971). Transitive inferences and memory in young children. *Nature*, 232(5311), 456-458.
- Buresh, J. S., & Woodward, A. L. (2007). Infants track action goals within and across agents. *Cognition*, 104, 287-314.
- Diesendruck, G., Salzer, S., Kushnir, T., & Xu, F. (2014). When choices aren't personal: The effect of statistical and social cues on children's inferences about the scope of preferences. *Journal of Cognition and Development*. Volume and page numbers??
- Fawcett, C., & Markson, L. (2010). Children reason about shared preferences. *Developmental Psychology*, 46(2), 299-309.
- Halford, G. S. (1984). Can young children integrate premises in transitivity and serial order tasks? *Cognitive Psychology*, 16, 65-93.
- Jern, A., Lucas, C. G., & Kemp, C. (2011). Evaluating the inverse decision-making approach to preference learning. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, & K. Weinberger (Eds.), *Advances in Neural Information Processing Systems 24* (pp. 2276-2284).
- Kushnir, T., Xu, F., & Wellman, H. (2010). Young children use statistical sampling to infer the preferences of others. *Psychological Science*, 21, 1134-1140.
- Lucas, C., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., Markson, L., & Hu, J.

- (2014.) The child as econometrician: A rational model of preference understanding in children. *PLOS One*, 9(3), e92160. doi:10.1371/journal.pone.0092160.
- Lumeng, J. C., Cardinal, T. M., Jankowski, M., Kaciroti, N., &Gelman, S. A. (2008). Children's use of adult testimony to guide food selection. *Appetite*, 51, 302-310.
- Luo, Y., &Baillargeon, R. (2007). Do 12.5-month-old infants consider what objects others can see when interpreting their actions? *Cognition*, 105, 489-512.
- Luo, Y., & Johnson, S. C. (2009). Recognizing the role of perception in action at 6 months. *Developmental Science*, 12, 142-149.
- Ma, L., & Xu, F. (2011). Young children's use of statistical sampling evidence to infer the subjectivity of preferences. *Cognition*, 120(3), 403-411.
- Mascaro, O., &Csibra, G. (2014). Human infants' learning of social structures: The case of dominance hierarchy. *Psychological Science*, 25(1), 250-255.
- Mou, Y., Province, J. M., &Luo, Y. (2014). Can infants make transitive inferences? *Cognitive Psychology*, 64, 98-112.
- Papafragou, A., &Musolino, J. (2004). Scalar implicatures: Experiments at the Semantics Pragmatics Interface. *Cognition*, 86(3), 253-282.
- Papafragou, A., &Tantalou, N. (2003). Children's computation of implicatures. *Language Acquisition*, 12(1), 17-82.
- Pears, R., & Bryant, P. (1990). Transitive inferences by young children about spatial position. *British Journal of Psychology*, 81(4), 497-510.
- Piaget, J. (1928). *Judgment and reasoning in the child*. London: Kegan Paul, Trench, &Trubner.
- Piaget, J. (1955). *The children's construction of reality*. London: Routeledge&Kegan Paul.
- Repacholi, B. M., &Gopnik, A. (1997). Early reasoning about desires: Evidence from 14- and

- 18-month-olds. *Developmental Psychology*, 33(1), 12-21.
- Robson, S. J., Lee, V., Kuhlmeier, V. A., & Rutherford, M. D. (2014). Infants use contextual contingency to guide their interpretation of others' goal-directed behavior. *Cognitive Development*, 31, 69-78.
- Schulz, L., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, 103(3), 322-332.
- Shultz, T. R., Pardo, S., & Altmann, E. (1982). Young children's use of transitive inference in causal chains. *British Journal of Psychology*, 73(2), 235-241.
- Smith, L. B., & Yu, C. (2008). Infants rapidly learn word-reference mappings via cross-situational statistics. *Cognition*, 106(3), 1558-1568.
- Sommerville, J. A., Woodward, A. L., & Needham, A. (2005). Action experience alters 3-month old infants' perception of others' actions. *Cognition*, 96, B1-B11.
- vonFersen, L., Wynne, C. D. L., Delius, J. D., & Staddon, J. E. R. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 17(3), 334-341.
- Vouloumanos, A., & Werker, J. F. (2009). Infants' learning of novel words in a stochastic environment. *Developmental Psychology*, 45(6), 1611-1617.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69, 1-34.
- Zentall, T. R., & Sherburne, L. M. (1994). Transfer of Value from S+ to S- in a Simultaneous Discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*, 20(2), 176-183.

Figure 1. Materials used in Experiment 1.

Figure 2. Children's responses in Experiment 1 test questions. In Test Question 1, children were asked whether Duckie would prefer object A, chosen in 100% of demonstration trials it appeared in, or object B, chosen in 70% of demonstration trials it appeared in. In Test Question 2, children were asked to infer Duckie's preferences for each of the three familiar objects seen in demonstration trials to a novel object.

Figure 3. Children's responses in Experiment 2 test questions. In Test Question 1, children were asked whether Duckie would prefer object A, chosen in 83% of demonstration trials it appeared in, or object B, chosen in 63% of demonstration trials it appeared in. In Test Question 2, children were asked to infer Duckie's preferences for each of the three familiar objects seen in demonstration trials to a novel object.

Table 1. Children's rate of selection of stimulus objects in baseline questions from Experiments 1 & 2.

Table 2. Experiment 1, question 2 response patterns of children who chose the 100% object over 70% object in question 1.

Table 3. Experiment 1, question 2 response patterns of all children in Experiment 1.

Table 4. Experiment 2, question 2 response patterns of children who chose the 83% object over 63% object in question 1.

Table 5. Experiment 2, question 2 response patterns of all children in Experiment 2.

Figure 1

Exp 1 Q1: 100% object v. 70% object

Exp 1 Q2: Comparisons to novel object

Figure 2

Exp 2 Q1: 83% object v. 63% object

Exp 2 Q2: Comparisons to novel object

Figure 3

Table 1

Children's rate of selection of stimulus objects in baseline questions from Experiments 1 & 2.

Object #	Object description	Number of times object appeared as baseline option	Number of times selected by a participant	Rate of selection
1	Candy dispenser	58	28	48%
2	Disc	64	28	44%
3	J-shaped object	62	33	53%
4	L-shaped object	58	36	62%

Table 2

Experiment 1, question 2 response patterns of children who chose the 100% object over 70% object in question 1.

Choice in 2a	Choice in 2b	Choice in 2c	Number of children
100	70	20	2
100	70	novel	11
100	novel	novel	5
novel	novel	novel	0
100	novel	20	3
novel	70	20	1
novel	70	novel	2
novel	novel	20	3

Note: Response patterns suggesting individual children inferred $A > B > C$ are bolded.

Table 3

Experiment 1, question 2 response patterns of all children in Experiment 1.

Choice in 2a	Choice in 2b	Choice in 2c	Number of children
100	70	20	4
100	70	novel	11
100	novel	novel	5
novel	novel	novel	0
100	novel	20	4
novel	70	20	1
novel	70	novel	3
novel	novel	20	3

Note: Response patterns suggesting individual children inferred $A > B > C$ are bolded.

Table 4

Experiment 2, question 2 response patterns of children who chose the 83% object over 63% object in question 1.

Choice in 2a	Choice in 2b	Choice in 2c	Number of children
83	63	29	5
83	63	novel	8
83	novel	novel	3
novel	novel	novel	4
83	novel	29	4
novel	63	29	0
novel	63	novel	0
novel	novel	29	0

Note: Response patterns suggesting individual children inferred $A > B > C$ are bolded.

Table 5

Experiment 2, question 2 response patterns of all children in Experiment 2.

Choice in 2a	Choice in 2b	Choice in 2c	Number of children
83	63	29	7
83	63	novel	8
83	novel	novel	6
novel	novel	novel	5
83	novel	29	4
novel	63	29	0
novel	63	novel	3
novel	novel	29	2

Note: Response patterns suggesting individual children inferred $A > B > C$ are bolded.