The development of heuristics in children: Base-rate neglect and representativeness

Samantha Gualtieri (sgualtieri@uwaterloo.ca) & Stephanie Denison (stephanie.denison@uwaterloo.ca)
University of Waterloo, Department of Psychology

Abstract

This paper examines the development of the representativeness heuristic in early childhood. Using a novel paradigm, we investigated 3- to 6-year-old children's ability to use base-rate and individuating information in their predictive inferences. In Experiment 1, we presented children with base-rate and individuating information separately to test their ability to use each independently. In Experiment 2, we presented children with base-rate and individuating information together. Two critical trial types were used, one in which the base-rate information and individuating information pointed to the same response and one in which the base-rate and individuating information pointed to conflicting responses. Results suggest that children progress to adult-like heuristic-based responding at 6 years of age.

Keywords: decision making, base-rate neglect, heuristics

Introduction

From their first year of life, human infants are able to use numerical information in their inductive inferences. For instance, infants expect a random sample drawn from a larger distribution to resemble the larger distribution (Teglas et al., 2007; Xu & Garcia, 2008). However, adults tend to under use numerical base-rate information and over rely on other types of information in their predictive inferences (Kahneman & Tversky, 1973). These two literatures pose a conundrum for cognitive developmentalists: if humans can use base-rates to make inferences as early as 12 months, why are adults so inclined to neglect them and when does the tendency to do so emerge?

In their seminal work, Kahneman and Tversky (1973) outlined notable cases of adults neglecting base-rates. In the classic lawyer-engineer problem, participants read a personality description that was randomly selected from a sample of 70 lawyers and 30 engineers. This description was of a person who was conservative, enjoyed puzzles, and did not care for social issues. When asked to report whether the person was a lawyer or an engineer, participants reported it was more likely that the person was an engineer. That is, people neglected base-rates (i.e., the number of lawyers and engineers) and over used individuating information (i.e., the personality description) in their estimates. Kahneman and Tversky (1973) referred to this as the representativeness heuristic: people over rely on individuating information that fits their representation of a group's characteristics and ignore important factors, such as base-rates, resulting in biased judgments in some cases.

The vast literature examining normative (often also referred to as analytic) versus non-normative (often also referred to as heuristic) responding often includes two assumptions (see Kokis et al., 2002 for a review). The first assumption, which has received mixed empirical support, is that heuristic responding should decrease in adult samples as intelligence increases (adults of higher intelligence should override a heuristic response more often, in favour of the normative response). A related, though largely untested assumption is that heuristic responding should decrease with age, as the ability to override the heuristic response should improve over development. However, it is difficult to draw conclusions from the few studies that have investigated the development of base-rate neglect and the representativeness heuristic in childhood, as their designs lack features that have proven vital to understanding the adult work. One pivotal aspect of Kahneman and Tversky's (1973) problems is the familiarity of the group information to participants. In the previous lawyer-engineer example, adults believe the person is an engineer, as the description they are given closely resembles their representation or stereotype of an engineer. In order to use this heuristic, one has to be familiar with how individuating information can be used to classify a specific case. It is important to consider young children's experience with group information when examining their use of the representativeness heuristic (Stanovich, West, & Toplak, 2011). For example, Jacobs and Potenza (1991) and Davidson (1995) found age-related *decreases* in base-rate use, with the youngest children (6- and 7-year-olds) providing more normative responses than older children and adults in tasks that were adapted for children. These results were interpreted as revealing a counter-intuitive idea: younger children actually make better predictive judgments than older children and adults, as they were more likely to provide the response that was closer to the base-rates. However, participants in these studies were presented with group information that may have been unfamiliar to the youngest children in the sample (e.g., cheerleader and band member stereotypes). Since it is reasonable to assume that the youngest children did not have the relevant group information to begin with, children's use of base-rates may have arose from being unaware of the category information, rather than overriding a heuristic response with a normative one (Stanovich et al., 2011).

One previous study has highlighted this important problem by attempting to manipulate the familiarity of stereotype information across ages. De Neys and Vanderputte (2011) investigated 5- and 8-year-old children's responses to base-rate problems that used both familiar and unfamiliar group information. Their results supported the claim that younger children might be providing normative responses only when they are unfamiliar with the presented stereotypes (and thus only have base-rates to go on). However, attempting to find ages at which these stereotypes are emerging would help determine if the representativeness heuristic is used when both types of information are available to the child.

Using a novel paradigm, we examined young children's use of base-rate and individuating information in their inferences. In Experiment 1, we established whether 3- to 5-year-old children could use base-rate and individuating information when presented alone. No previous studies have established whether children could use base-rate and individuating information separately in their tasks, which is vital to understanding how children consider this information when presented together. In Experiment 2, we presented 3- to 6-year-old children with both base-rate and individuating information. Analogous to the adult literature, we investigated children's inferences when presented with individuating information that conflicted with base-rate information. To examine the integration of information in young children further, we also included problems in which the two types of information did not conflict.

Experiment 1: Methods

Participants

Children in both experiments were tested individually in lab, at their school or daycare, or at a local museum. Forty-eight children participated in the base-rate condition, including 16 three-year-olds, 16 four-year-olds, and 16 five-year-olds (*Mean age* = 4 years, 4 months, *range* = 3 years, 0 months, 9 days to 5 years, 11 months, 20 days, *female* = 20). An additional child was tested but excluded for failing to correctly identify the majority group on both problems (see Procedure for details). Forty-eight children participated in the individuating condition, including 16 three-year-olds, 16 four-year-olds, and 16 five-year-olds (*Mean age* = 4 years, 6 months, *range* = 3 years, 0 months, 27 days to 5 years, 11 months, 29 days, *female* = 24). An additional child was tested but excluded for non-compliance.

Procedure, Design and Predictions

Participants in both conditions completed two problems, presented consecutively. The experimental session was presented to the children on a laptop, using a PowerPoint presentation, which was narrated live by an experimenter. On the first slide, children were told that they were going to hear about some robots on another planet, and that they would have to answer some questions about the robots after.

In the base-rate condition, each participant completed a color and a shape problem, in counterbalanced order (see Figure 1 for a diagram of the base-rate condition procedure). Two problem-types (color and shape) were used to maintain children's attention. In the color problem, following the initial introduction, participants next saw a slide with two robots standing side by side at a library, one robot wearing red and another robot wearing purple. On the following two slides, participants saw each robot by itself, and the researcher pointed out what color the robot was wearing. The next slide included both of the robots standing side by side again, and the researcher asked the child to point to the one wearing red, and the one wearing purple. On the next slide, participants saw a group of ten robots at the library, eight wearing red, and two wearing purple. Children were asked to indicate which type of robot there was more of. After the child provided their answer, the experimenter indicated that there were lots of robots wearing red, and very few wearing purple. Following this, children saw a slide with a single robot wearing a white coat, making its type unclear. Children were asked to recall which type of robot there was more of in the group. Participants were then asked to indicate which type of robot they thought the one wearing the white coat was. This procedure was also followed for the shape problem; however, participants saw one type of robot wearing hearts and another type of robot wearing stars. To ensure participants did not confuse the robots with those from the previous problem, the shape problem took place at a grocery store.

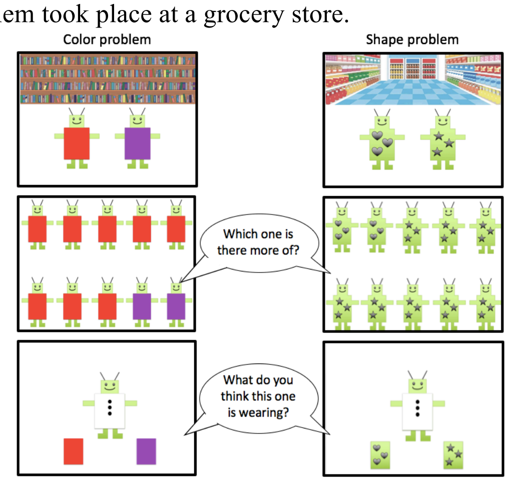

Figure 1: Base-rate condition procedure.

In the individuating condition, each participant completed a trait and a gender problem in counterbalanced order (see Figure 2 for a diagram of the individuating condition procedure). Young children have successfully considered gender stereotypes and trait-based categories in other work (Heyman & Gelman, 2000; Martin & Ruble, 2004). As adults are familiar with occupational stereotypes prior to completing the lawyer-engineer problem, we chose both of these trial types based on young children's familiarity with these categories. In the trait problem, following the initial introduction, participants saw a slide with two robots standing side by side at a park, one robot wearing blue and another robot wearing green. They were told that most of the time the robots wearing blue were nice, and the ones wearing green were naughty. On the next slide, the robot wearing blue was holding flowers, and participants were told it would bring flowers to another robot on its birthday. On the following slide, the robot wearing green was hiding presents, and participants were told that this robot would hide another robot's birthday presents. Participants then saw the blue and green robot standing side by side on the following slide. Participants were asked to indicate which type of robot was nice most of the time, and which type of robot was naughty most of the time. Following this, children saw a slide with a single robot wearing a white coat, making its type unclear. Children were asked to recall which traits were associated with each color. The experimenter then provided some additional information about the robot wearing the white coat, which matched one of the traits. Half of the children heard a description of nice behavior, as the experimenter said that it liked to clean up the park and help its friend play on the swing. The other half of the children heard a description of naughty behavior, as the experimenter said that it liked to make a mess of the park and scare other robots. Importantly, there was no mention of either the word "naughty" or "nice"; children were only told about the mystery robot's behaviour. Participants were asked to indicate which type of robot they thought the one wearing the white coat was. This procedure was also followed for the gender problem, however participants saw one type of robot wearing yellow and another type of robot wearing orange. They were told that most of the time the robots wearing yellow liked to play with toys that girls like, and those wearing orange liked to play with toys that boys like. When introduced to the yellow robot individually, it had mostly toys typically associated with girls (i.e., a unicorn and a doll). When introduced to the orange robot individually, it had mostly toys typically associated with boys (i.e., a dinosaur and a helicopter). The individuating information was as follows: Half of the children heard a description more typical of girls, as the experimenter said it liked to play dress-up and house. The other half of the children heard a description more typical of boys, as the experimenter said that it liked to play with trucks and train sets. To ensure participants did not confuse the robots with those from the previous problem, the gender problem took place at a school.

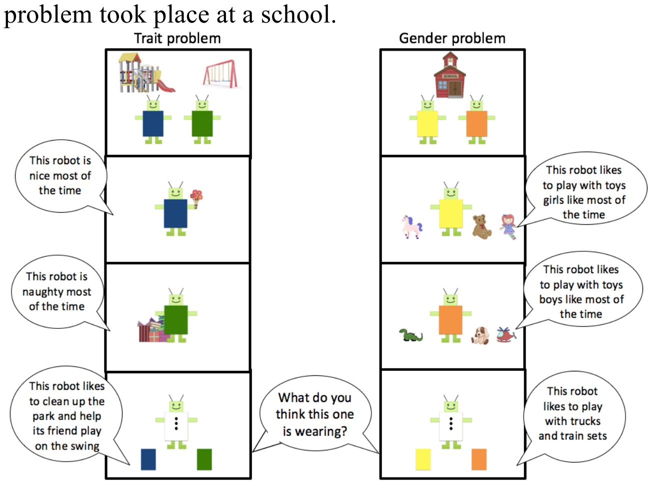

Figure 2: Individuating condition procedure.

Design In both conditions, the position of the two types of robots was counterbalanced in each problem (e.g., half of the children saw the orange robot on the left and half saw it on the right when they were introduced). The experimenter always began by telling the child about the robot on the left. In the base-rate condition, the robots that comprised the majority group were counterbalanced across both problems. The position of the majority group was also counterbalanced in both problems, as the majority was either presented as the first eight or last eight robots in the group of ten. In the individuating condition, the individuating information given about the robot wearing the white coat was counterbalanced as representative of each type.

Predictions If children can use base-rate information to make a predictive inference, then they should choose the type of robot that corresponds to the majority group in that condition. If children can use individuating information (specifically, the individuating information about these trait-based categories and gender stereotypes) to make a predictive inference, then they should choose the color corresponding to the individuating information given in that condition. Thus, children received a score of 1 for a correct response in each condition on each problem.

Experiment 1: Results

Base-rate condition. A repeated-measures Analysis of Variance (ANOVA) with score on each problem (color, shape) as a within-subjects factor and age (3-, 4-, 5-year-olds) as a between-subjects factor revealed a significant interaction (F (2, 45) = 7.82, p = .001, η^2_p = .26). No main effects of age (p = .11), or problem type (p = .43) were found. Post hoc analyses revealed that the interaction was driven by differences in performance on the shape problems, as 3-year-olds significantly differed from 4-year-olds ($Mean_{Difference}$ = -.44, p = .01) and 5-year-olds ($Mean_{Difference}$ = -.44, p = .01).

Overall performance was significantly different from chance, as children tended to indicate that the robot belonged to the majority group (M = 1.5, SD = .77, t (47) = 4.49, p < .001). As we were interested in development across age, we also examined children's performance at each age separately. T-tests revealed that the performance of 4-year-old (M = 1.56, SD = .73, t (15) = 3.09, p = .007) and 5-year-old (M = 1.75, SD = .68, t (15) = 4.39, p = .001) children were significantly different from chance. Three-year-olds' overall performance did not differ from chance (M = 1.19, SD = .83, p = .34). To investigate this further, we compared 3-year-olds' performance on color versus shape problems with a paired-samples t-test to determine whether they found one problem more difficult than the other. There was a significant difference in performance (t (12) = 2.31, p = .040), as 3-year-olds performed above chance on color problems (M = .75, SD = .45, t (16) = 2.24, p = .041), but not on shape problems (M = .44, SD = .51, p = .63).

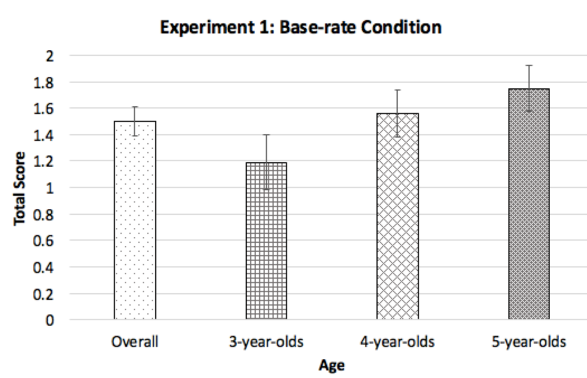

Figure 3: Base-rate condition results by age.

Individuating condition. A repeated-measures ANOVA with score on each problem (gender, trait) as a within-subjects factor and age (3-, 4-, 5-year-olds) as a between-subjects factor revealed a main effect of age on performance (F (2, 45) = 7.12, p = .002, η^2_p = .24), and no main effect of problem type (p = 1), or interaction (p = .44). Post-hoc tests revealed the main effect of age was driven by differences in the performance of 3- and 5-year-old children ($Mean_{Difference}$ = .88, p = .001).

Overall performance was significantly different from chance, as children indicated that the robot belonged to the group the individuating information corresponded with (M = 1.42, SD = .74, t (47) = 3.91, p < .001). Additional t-tests revealed 5-year-old children performed significantly above chance (M = 1.88, t (15) = 10.25, p < .001). However, 4-year-olds were only marginally different from chance (M = 1.38, p = .08), and 3-year-olds did not exceed chance (M = 1, p = 1).

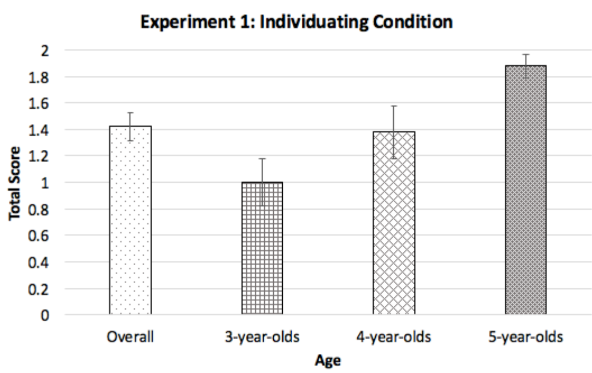

Figure 4: Individuating condition results by age.

Experiment 1: Discussion

Using a novel paradigm, we established that children are able to use base-rate and individuating information separately to make a predictive inference, but that these abilities are still developing between 3 and 5 years. When base-rate information was made salient, 4- and 5-year-old children used base-rates to predict group membership at above chance levels. Although 3-year-old children's overall performance was not above chance, they were able to use base-rate information when group membership was presented using color. When individuating information was made salient, 5-year-old children used the individuating information to predict group membership at above chance levels. Four-year-olds approached significance, using the information correctly approximately 70% of the time. However, 3-year-old children performed at chance levels, suggesting that they are unable to use individuating information to make a predictive inference in our task.

Having examined the ability to use each type of information separately in young children, the next step is to extend the task to an analogue of the classic lawyer-engineer problem in Experiment 2. We continued to test 3-, 4- and 5-year-olds in Experiment 2, as these age groups displayed a range of abilities for using the individuating information presented to them. Each age group successfully used base-rate information when presented using color, so all problems in Experiment 2 use color. This set-up, in which all ages can rely on base-rates, thus holding that ability more or less constant but varying the ability to use individuating information, could be particularly informative. It allows investigation of questions such as: At what ages are children able to integrate information from both of these sources? Will children show bias based on representativeness as soon as they can reliably use individuating information (at 5 years of age)? Or does this bias develop over time (will there be a delay in its appearance)? To examine this progression, the age range is extended to 6 years.

Experiment 2: Methods

Participants

Sixty-four children participated, including 16 three-year-olds, 16 four-year-olds, 16 five-year-olds, and 16 six-year-olds (*Mean age* = 5 years, 1 month, *range* = 3 years, 6 months, 24 days to 6 years, 11 months, 27 days, *female* = 31). An additional seven children were tested, but were excluded due to interference from another child (n = 3) or failing to correctly identify the majority group in both problems (n = 4).

Procedure, Design and Predictions

Participants completed two problems, presented consecutively. As in Experiment 1, children completed a trait and a gender problem, using the same color, group, and individuating information as before. Each child saw one color pair and the associated group information on the first trial, and another color pair and associated group information on the second trial.

Using the trait problem as an example, children were first introduced to the two types of robots, and were told about the traits that corresponded to each color, just as in Experiment 1 (e.g., nice blue and naughty green robots). Once again, they heard some information about the robots' behavior at a birthday party. Following this, participants saw the robots side by side and were asked to indicate which robot was nice and which one was naughty most of the time. Next, participants saw a group of ten robots, with eight wearing blue, and two wearing green. Children were asked to indicate which type of robot there was more of at the park. After the child provided their answer, the experimenter indicated that there were lots of robots wearing blue, and just a few wearing green. Following this, children saw a slide with a single robot wearing a white coat, making its type unclear. To ensure both pieces of information were equally salient to the child, the experimenter reminded the child in a counterbalanced order that most of the time the blue robots were nice and the green robots were naughty, and that there were more robots wearing blue at the park. In each problem, base-rate and group information were explicitly mentioned twice to the child prior to hearing the individuating information about the robot. The experimenter then provided some additional information about the robot wearing the white coat, which corresponded to one of the group traits. As in Experiment 1, for the "nice" trait, the experimenter said that it liked to clean up the park and help its friend play on the swing. For the "naughty" trait, the experimenter said that it liked to make a mess of the park and scare other robots. Participants were asked to indicate which type of robot they thought the one wearing the white coat was.

Each participant completed a conflict and a no-conflict problem. In the conflict problem, the base-rate conflicted with the individuating information given about the unknown robot. For instance, if there were eight nice robots, the unknown robot was described as liking to make a mess of the park and scare other robots. In the no-conflict problem, the base-rate and individuating information both cued the same group. For instance, if there were eight nice robots, the unknown robot was described as liking to clean up the park and help its friend play on the swing.

Design The order in which the child completed the conflict and no-conflict problems was counterbalanced across participants. Further, the group information associated with the robots in the conflict and no-conflict problems was counterbalanced (i.e., half of the children completed a conflict problem about the trait story, and the other half completed a no-conflict problem about the trait story). The order of reminding about base-rate and group information was also counterbalanced. All factors that were counterbalanced in Experiment 1 were also counterbalanced here (i.e., color and position of majority, the position of the two types of robots, the type of robot the individuating information represented).

Predictions In both problems, the normatively correct answer is to produce a response that aligns with the base-rates (and in the no-conflict problems, this also corresponds with a response from individuating information). In no-conflict problems, children should provide the response consistent with both the base-rate and individuating information, as these will point to the same response. In conflict problems, if children are biased toward using individuating information when it conflicts with the base-rate, they should choose the color corresponding to the individuating information, rather than the base-rate information, similarly to adults in the classic tasks. However, if children do not yet have this reasoning bias, or if they do not fully grasp the stereotype/category information, then they might provide responses consistent with base-rates, rather than individuating information. As each child completed one conflict and one no-conflict problem, they received a score of 0 or 1 for each problem type. In both problems, children received a score of 1 if they chose the majority group.

Experiment 2: Results

A repeated-measures ANOVA including score on each problem type (conflict, no-conflict) as a within-subjects factor and age (3-, 4-, 5-, 6-year-olds) and information recapped first (base-rate, individuating) as between-subjects factors revealed a significant interaction of age and problem type (F (3, 56) = 5.66, p = .002, η^2_p = .23). There was a main effect of problem type (F (1, 56) = 18.67, p < .001, η^2_p = .25). There were no other main effects or interactions.

The interaction was driven by an effect of age on the no-conflict problems ($Mean_{No\text{-}conflict}$ = .70, SD = .46; F (3, 60) = 5.91, p = 0.001). Post-hoc analyses revealed 6-year-olds' scores on no-conflict problems significantly differed from 3-year-olds ($Mean_{Difference}$ = .44, p = .02) and 4-year-olds ($Mean_{Difference}$ = .56, p = .002). The difference between 4- and 5-year-old children on no-conflict problems also approached significance ($Mean_{Difference}$ = -.38, p = .06; see Figure 5). There was no effect of age on performance in the conflict problems ($Mean_{Conflict}$ = .38, SD = .49; p = .38).

Additional t-tests revealed age differences on conflict and no-conflict problems. On conflict problems, 3-year-olds (M = .44, SD = .51, p = .63), 4-year-olds (M = .44, SD = .51, p = .63), and 5-year-olds (M = .44, SD = .51, p = .63) did not exceed chance. Six-year-old children were significantly different from chance (M = .19, SD = .4, p = .007). On no-conflict problems, 3-year-olds (M = .56, SD = .51, p = .63) and 4-year-olds (M = .44, SD = .51, p = .63) did not exceed chance. Five-year-olds were significantly different from chance (M = .81, SD = .40, t (15) = 3.10, p = .007), and six-year-olds were at ceiling (M = 1).

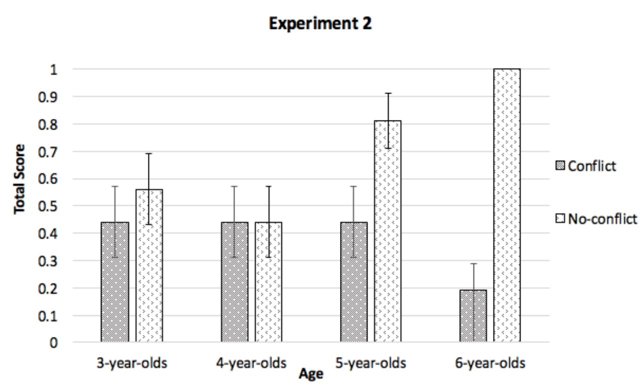

Figure 5: Experiment 2 results by age.

Experiment 2: Discussion

The results of Experiment 2 revealed interesting age differences in children's use of base-rates. Three- and four-year-olds performed at chance levels on both no-conflict and conflict problems, even though the normative response in the no-conflict problems corresponded with both base-rate and individuating information. Perhaps the 4-year-olds, who were starting to use individuating information on its own, but not entirely reliably, face issues integrating both pieces of information at first, due to a cognitive decoupling problem (see Stanovich, et al., 2011). The responses of 5-

and 6-year-old children in our task most strongly resembled the typical patterns of adults[1]. In no-conflict problems, both 5- and 6-year-old children provided the normative response at above chance levels, suggesting that by 5, children can fully handle the task. Interestingly, 5-year-old children did not show a pronounced bias for relying on individuating information, as they were at chance on conflict problems. Six-year-olds showed the predicted pattern from adults on the conflict problems, suggesting that this bias is fairly well ingrained, even at this young age.

General Discussion

In two experiments, we were able to verify methods for examining base-rate neglect and the representativeness heuristic in young children. We established the ages at which children can use base-rate and individuating information, when presented separately and made highly salient, in Experiment 1. Based on the extent to which children were able to use base-rate and individuating information in Experiment 1, we found interesting age differences in performance when both pieces of information were available in Experiment 2. Although three-year-olds could use base-rate information independently, they did not show a preference for base-rate information on conflict or no-conflict problems. Four-year-olds, who were able to use both base-rate and individuating information independently, also struggled when both pieces of information were presented together. Five-year-old children performed above chance on no-conflict problems, though they did not consistently rely on base-rate or individuating information on conflict problems. By the age of 6, children's responses revealed a pattern reminiscent of adult performance. These age differences provide useful insight on the representativeness heuristic, suggesting that its emergence may not be as straightforward as previously thought. Further research with 4- and 5-year-old children would provide insight on the nuances in the development of this bias.

Much additional research can be conducted to further examine the development of base-rate neglect. For example, ongoing work in our lab is investigating a second important feature of the adult work on representativeness that has been, to our knowledge, untested in children. That is, of course in many cases, it is reasonable (and highly computationally efficient) to rely on the type of individuating information that is provided in these vignettes, as the information appears to be accurate and reliable. However, adults will neglect base-rates in favor of individuating information even when that individuating information is presented as unreliable or uninformative for classification. This suggests that adults have an overwhelming expectation that social information should always be heavily relied on. Have 6-year-old children developed this expectation? Or will they be more sensitive to factors like reliability as this bias is developing? Current experiments in our lab are examining this and related questions that should provide further insight into the development of reasoning biases.

[1] Adults ($N = 36$) completed an experiment identical to Expt. 2 (i.e., same gender and trait conflict and no-conflict problems). Space does not permit a full report of the experiment, but adults produced the expected results (giving the normative response for no-conflict and the non-normative response for conflict problems).

Acknowledgments

We thank the families for participating and Emily Chan and Elizabeth Attisano, for their assistance with data collection.

References

Davidson, D. (1995). The representativeness heuristic and the conjunction fallacy effect in children's decision making. *Merrill Palmer Quart, 41,* 328–346.

Heyman, G.D., & Gelman, S.A. (2000). Preschool children's use of trait labels to make inductive inferences. *J Exp Child Psychol, 77*, 1 – 19.

Jacobs, J.E., & Potenza, M. (1991). The use of judgment heuristics to make social and object decisions: A developmental perspective. *Child Dev, 62*, 166–178.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*(4), 237–251.

Kokis, J., Macpherson, R., Toplak, M., West, R.F., & Stanovich, K.E. (2002). Heuristic and analytic processing: Age trends and associations with cognitive ability and cognitive styles. *J Exp Child Psychol, 83,* 26–52.

Martin, C.L., & Ruble, D. (2004). Children's search for gender cues. *Curr Dir Psychol Sci, 13*, 67 – 70.

Stanovich, K. E., West, R. F., & Toplak, M. E. (2011). The complexity of developmental predictions from dual process models. *Dev Rev, 31*, 103-118.

Teglas, E., Girotto, V., Gonzalez, M., & Bonatti, L. L. (2007). Intuitions of probabilities shape expectations about the future at 12 months and beyond. *P Natl Acad Sci USA, 104,* 19156–19159.

Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *P Natl Acad Sci USA, 105,* 5012–5015.