

Learning By Asking: How Children Ask Questions To Achieve Efficient Search

Azzurra Ruggeri (a.ruggeri@berkeley.edu)

Department of Psychology, University of California, Berkeley, USA
Max Planck Institute for Human Development, Berlin, Germany

Tania Lombrozo (lombrozo@berkeley.edu)

Department of Psychology, University of California, Berkeley, USA

Abstract

One way to learn about the world is by asking questions. We investigate how children ($n=287$, 7- to 11-year olds) and young adults ($n=160$ 17- to 18-year olds) ask questions to identify the cause of an event. We find a developmental shift in children's reliance on hypothesis-scanning questions (which test hypotheses directly) versus constraint-seeking questions (which reduce the space of hypotheses), but also that all age groups ask more constraint-seeking questions when hypothesis-scanning questions are unlikely to pay off: when the problem is difficult (Studies 1 and 2) or the solution is one among equally likely alternatives (Study 2). These findings are the first to demonstrate that even young children adapt their strategies for inquiry to increase the efficiency of information search.

Keywords: information search, active learning, 20-questions game, cognitive development.

Introduction

Yesterday, a man was late for work. Why?

In everyday life, we constantly encounter events that require explanation. An observation, such as a man being late for work, can prompt us to seek underlying causes, either out of curiosity or because our future actions depend on them. For example, if the man was late because there is a public transportation strike, we might do well to rethink our own travel plans.

Although the process of generating and evaluating explanatory hypotheses plays a crucial role in learning and development (Lombrozo, 2006), little is known about how children test hypotheses by seeking novel information from knowledgeable informants, or how this ability changes over the course of development. In particular, how do children ask questions to arrive at the correct hypothesis?

In this paper we investigate children's ability to flexibly adapt the kinds of questions that they ask, and whether the way in which they do so is responsive to task characteristics that affect the efficiency of different strategies. Specifically, we investigate how children (7- to 8-year-olds and 9- to 11-year-olds) and young adults (17- to 18-year-olds) identify the cause of an event by asking yes-or-no questions. Building on a classic study by Mosher and Hornsby (1966), which used a task like the game of "20-questions," we differentiate between *constraint-seeking* questions and *hypothesis-scanning* questions. Constraint-seeking questions attempt to reduce the space of the possible solutions by asking about features that could apply to multiple solutions

(e.g., "Was the man late because of something related to his means of transport?" or "Did something happen at home?"). In contrast, hypothesis-scanning questions are tentative solutions (e.g., "Was he late because he missed the bus?" or "Was he late because he overslept?"). Mosher and Hornsby found a large developmental change from age 6 to age 11, which has since been replicated (e.g., Denney, 1975; Denney & Denney, 1973), with younger children overwhelmingly asking hypothesis-scanning questions, and constraint-seeking questions becoming more frequent in the course of development.

Subsequent research using the 20-questions paradigm has found that different task features can influence children's and adults' reliance on constraint-seeking questions. For example, Siegler (1977) found that 13- and 14-year-old adolescents were strongly influenced by the order in which two isomorphic 20-questions problems were presented (see also Nelson, Divjak, Martignon, Gudmundsdottir, & Meder, 2013), and Ruggeri and Feufel (2013) found that describing objects at a basic level (e.g., "dog" as opposed to "retriever") increased the proportion of constraint-seeking questions in all age groups, suggesting that the basic-level representations facilitated children's ability to identify object-general features on which to base their questions. However, few studies have investigated children's ability to adapt their search for information as a function of task features that influence the efficacy—as opposed to the ease of implementation—of different questions. One exception is Nelson et al. (2013), who found that 4th grade children were reasonably effective at asking constraint-seeking questions that partitioned a space of alternatives into equal subgroups, especially when the statistics of the possible solutions matched those of the real world.

While several factors seem to influence children's general ability to ask constraint-seeking questions, it remains an open question whether their strategies for inquiry are responsive to the anticipated efficiency of a given *strategy* (i.e., constraint seeking versus hypothesis scanning) in the task at hand. One possibility is that children's strategies for inquiry, compared to those of young adults, are quite inflexible. If this is the case, young children will overwhelmingly rely on hypothesis-scanning questions and will do so regardless of how effective the strategy is for a given instantiation of the task. This possibility is broadly consistent with the findings to date that have used the 20-questions paradigm: hypothesis scanning appears to robustly

dominate in early childhood, with shifts towards constraint seeking induced by factors that reduce the strategy's cognitive demands rather than those that increase its efficiency relative to hypothesis scanning.

Another possibility is that, despite a general tendency to rely on hypothesis-scanning questions, children will appropriately modify their strategies for inquiry, engaging in less hypothesis scanning when it's less likely to pay off. Such flexibility is broadly consistent with previous research on how children seek information in the process of decision-making, such as deciding which story a fictitious character would like the most (Davidson, 1991b). Moreover, even young children's predecisional information search is influenced by some task characteristics, such as the amount of information available (Davidson, 1991a), and by age 10 children become more likely to use relevant information to guide their search (Davidson, 1991b).

To investigate the flexibility of children's strategies for inquiry, we examined whether children are more inclined toward constraint-seeking questions when encountering an unfamiliar domain (Study 1), a difficult problem (Studies 1 and 2), or a set of possibilities that are equally likely in light of prior beliefs (Study 2). Overall, we predicted that both children's and young adults' questions would be responsive to these task characteristics. Specifically, we tested the following three predictions. First, familiarity with a scenario suggests richer prior knowledge about the possible causes of the situation described. For example, one might know that it is more likely that a man would be late for work because he was caught in a traffic jam than because his house was flooded during the night. This background knowledge could, in turn, influence the attractiveness of a hypothesis-scanning strategy that first tests the most likely hypotheses, because this strategy has the reasonable potential for a "quick win." We therefore hypothesized, in Study 1, that participants would be more likely to adopt a hypothesis-scanning strategy for scenarios that were familiar to them (e.g., a man being late for work, or a boy being late for school) than for those that were unfamiliar (e.g., an alien arriving late for a reunion).

Second, in both Study 1 and Study 2 we varied the prior probability of the task solution: it was either unlikely (e.g., the man was late because he had to wait for a plumber to repair a water leak) or likely (e.g., the man was late due to traffic). We expected that children might be more likely to adopt a constraint-seeking strategy after a hypothesis-scanning strategy failed to deliver a "quick win," and that this would occur more often when the solution designated as correct was a priori unlikely than a priori likely.

Third, we hypothesized that children (in Study 2) would be more likely to adopt a constraint-seeking strategy when trying to differentiate between candidate hypotheses that were all equally likely (such that a hypothesis-scanning approach would require selecting the hypotheses to test arbitrarily) than when aiming to differentiate hypotheses that varied in probability (such that a hypothesis-scanning strategy could begin with the most likely options). Finally,

from a more developmental perspective and in keeping with previous results (e.g., Mosher & Hornsby, 1966), we hypothesized a general linear developmental improvement in participants' performances, with an increase in the efficiency of the question-asking strategies adopted and a corresponding decrease in the number of questions needed to reach the solution.

Study 1

Method

Participants Participants were 96 children in second or third grade (66 female, $M_{\text{age}} = 7.51$ years; $SD = .50$), 87 children in fifth grade (56 female, $M_{\text{age}} = 9.83$ years; $SD = .77$), and 90 young adults (35 female, $M_{\text{age}} = 17.62$ years; $SD = 1.07$) from three schools in Livorno, Italy. The students represented a variety of social classes.

Design and procedure The experiment consisted of individual interviews. At the beginning the experimenter read the participant the task instructions, ensuring they were completely understood. Participants were presented with a short description of an event (e.g., "Yesterday, a boy was late for school") and asked why it occurred ("Why?"), after which they were expected to ask questions. Participants were randomly assigned to one of six possible experimental conditions, a 3×2 matrix resulting from the cross between two independent between-subjects variables, scenario and solution type (likely, unlikely), which we explain below.

Scenario. There were three possible scenarios: (1) "Yesterday, a man was late for work"; (2) "Yesterday, a boy was late for school"; and (3) "Yesterday, an alien was late for the supreme reunion." These scenarios were designed to vary in familiarity, with Scenario 1 being more familiar to young adults than to children, Scenario 2 familiar to all participants, and Scenario 3 unfamiliar to all participants.

Solution. There were two possible solutions for each scenario: (a) a *likely* solution, and (b) an *unlikely* solution (where anticipated perceived probabilities were confirmed on a post-test). Each solution was structured into three levels of causal detail. In the likely solution, the car/spacecraft taken by the man/boy/alien to go to work/school/the reunion was caught in a traffic jam (level 1) due to a car accident (level 2) that was caused by a driver who ran a red light (level 3). In the unlikely solution, the man/boy's father/alien had to wait for the plumber (level 1), whom he had called because the house flooded during the night (level 2) because a pipe had broken (level 3).

Results

Success rate In total, 105 participants did not finish the game: 70 younger children, 27 older children, and 8 young adults.

Number of questions needed to reach the solution For this analysis we considered only those participants who reached the complete solution (all three levels). We

performed a univariate analysis of variance (ANOVA) with the number of questions needed to reach the solution as the dependent variable and age group, scenario, and solution type as independent variables. This analysis revealed a main effect of age group, $F(2,167) = 4.22, p = .016, \eta^2 = .05$. A Bonferroni post hoc analysis revealed that older children asked more questions prior to reaching the complete solution ($M_{old_child} = 19.73, SD = 14.82$) than did either younger children ($M_{young_child} = 11.58, SD = 9.11, p < .001$) or young adults ($M_{young_adults} = 15.74, SD = 9.48, p = .029$). We found no difference between younger children and young adults ($p = .122$). We also found a main effect of solution type, $F(1,167) = 60.65, p < .001, \eta^2 = .29$: Participants assigned to the version of the game with the unlikely solution needed more questions to reach the solution than those who were assigned the likely solution. There were no additional significant effects.

Type of questions We coded questions as either hypothesis scanning or constraint seeking. The experimenter, an Italian student assistant blind to the experimental hypotheses, wrote down all the questions asked during the experiment. In addition, the experimental session was audio recorded, and based on notes and recordings, the experimenter coded all questions immediately after the session was over. All questions were additionally and independently coded following the session by a second Italian student assistant, blind to the experimental hypotheses, resulting in total agreement of Kappa = .993, $p < .001$. In the few cases where the two raters did not agree, a third Italian rater, blind to the experimental hypotheses and procedure, was consulted.

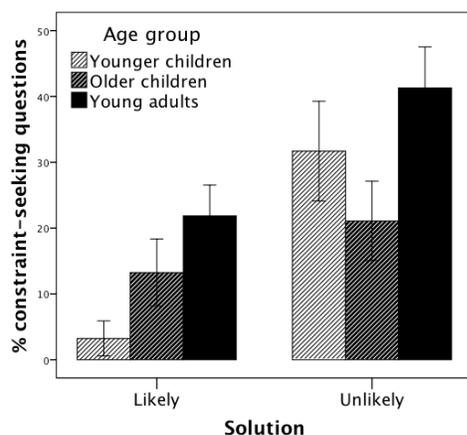


Figure 1. Study 1: Percentage of constraint-seeking questions, displayed by age group and solution type. Error bars represent one SEM in each direction.

For participants who reached the complete solution, we calculated the percentage of total questions that were constraint seeking (Figure 1). This percentage was analyzed as the dependent variable in a univariate ANOVA with age group, scenario, and solution type as independent variables. The analysis revealed a main effect of age, $F(2,167) =$

$17.08, p < .001, \eta^2 = .18$. A Bonferroni post hoc analysis confirmed that young adults asked a higher proportion of constraint-seeking questions ($M_{young_adult} = 30\%, SD = 19\%$) than older children ($M_{old_child} = 16\%, SD = 15\%, p < .001$), who in turn asked a higher proportion of constraint-seeking questions than younger children ($M_{young_child} = 5\%, SD = 10\%, p = .009$). We also found a main effect of solution type, $F(1,167) = 19.86, p < .001, \eta^2 = .12$. Participants asked a higher proportion of constraint-seeking questions in the game with an unlikely solution ($M_{unlikely} = 33\%, SD = 19\%$) compared to the game with a likely solution ($M_{likely} = 15\%, SD = 16\%$). Notably, we found no main effect of scenario on the type of questions asked, nor an interaction between age and solution.

First question To analyze the type of question that was asked first—hypothesis scanning versus constraint seeking—we conducted a logistic regression using age group, scenario, and solution type as predictors. For this analysis we included all participants, even those who did not succeed in finishing the game. The Wald criterion demonstrated that only age group ($p < .001$) made a significant contribution to predicting initial question type, whereas scenario and solution type were not significant predictors. The exp(B) value indicates that older age groups (older children compared to younger children, and young adults compared to older children) had a decreased likelihood of generating an initial hypothesis-scanning question (by .22 times), after controlling for the other factors in the model.

Study 2

Study 1 found that the proportion of constraint-seeking questions asked increased with age, replicating prior results, and additionally obtained the new result that the proportion of constraint-seeking questions increased when the solution was unlikely. Surprisingly, however, we did not find an effect of scenario, suggesting that familiarity did not influence participants' performances or search for information. However, this could be because participants judged some hypotheses to be more plausible than others, even in the unfamiliar cases. Study 2 revisits the question of whether prior knowledge about which hypotheses are more likely influences reliance on constraint-seeking versus hypothesis-scanning questions and does so with a more controlled experimental manipulation.

Method

Participants Participants were 58 children in second or third grade (30 female, $M_{age} = 7.0$ years; $SD = .59$), 46 children in fifth grade (23 female, $M_{age} = 10.2$ years; $SD = .73$), and 70 young adults (39 female, $M_{age} = 17.5$ years; $SD = .79$) from two schools in Livorno, Italy. The students belonged to various social classes.

Design and procedure Like Study 1, Study 2 consisted of individual interviews. The instructions presented to the

participants by the experimenter were identical to those used in Study 1. After being read the instructions, the participants were presented with the following situation: “Yesterday, a man was late for work. Why? The solution is one of the following.” The experimenter then took out 10 cards. On each card was displayed a different hypothesis, as well as its probability, both as a label (i.e., high, moderate or low probability) and in natural frequencies (10, 4 or 2 out of 40 times). The experimenter read each card aloud, in random order, then put it down on a table. The cards were left on the table until the end of the session. Participants were told that the correct solution was among these 10 and to begin asking yes-or-no questions to find out which one it was.

Participants were randomly assigned to one of four experimental conditions in a 2×2 design that crossed two independent variables: hypothesis distribution (uniform, mixed) and solution type (likely, unlikely).

Distribution. In the two *uniform distribution* conditions, the alternative hypotheses (i.e., possible solutions) provided to participants were designed to appear equally likely. In these conditions, the experimenter explicitly told participants that “all the alternatives are equally likely to be the correct solution.” In the two *mixed distribution* conditions, the hypotheses were designed such that two would be judged very likely to happen, four moderately likely to happen, and four very unlikely to happen. In these conditions, the experimenter presented the probability of each hypothesis in a natural frequency format: “Out of 40 times a man is late, 10 times [very likely]/4 times [moderately likely]/1 time [very unlikely] it is because...”

To select hypotheses that would be perceived as equally likely (in the uniform conditions) or very likely/moderately likely/very unlikely (in the mixed conditions), we pretested 20 statements on an independent sample of 25 adults. Participants in the pretest were asked to rate the probability of the described 20 events on a 10-point scale, from 0 (extremely unlikely) to 10 (extremely likely). Using these data, we were able to identify five statements that were judged very likely, five that were judged very unlikely, and two that were judged moderately likely. For each pretested statement, we constructed a “matched” item that was similar but distinct. For example, for the statement “He wasn’t feeling well when he woke up,” which was judged very unlikely, we constructed a second statement, “He hadn’t felt well during the night.” This allowed us to increase the total number of statements for the experiment and also ensured that pairs of statements involved common features that could provide a basis for asking constraint-seeking questions (e.g., “Did he not feel well at some point?”).

Solution. In the two *likely* conditions, the correct solution was very likely and was the same solution for both the uniform and mixed distribution conditions. For the *uniform/likely* condition, this meant that all of the candidate solutions provided were equally likely. For the *mixed/likely* condition, the candidate solutions varied in probability. In the two *unlikely* conditions, the correct solution was unlikely, and was the same solution for both the uniform

and mixed distribution cases. For the *uniform/unlikely* condition, this meant that all of the candidate solutions provided were equally unlikely. For the *mixed/unlikely* condition, the candidate solutions varied in probability and were identical to those in the *mixed/likely* condition.

Results

Success rate All participants completed the task. Unlike Study 1, Study 2 involved 10 pre-specified candidate solutions; even a hypothesis-scanning strategy that involved selecting hypotheses to test at random would reach the solution with a maximum of 10 questions.

Number of questions needed to reach the solution We analyzed the number of questions required to reach the complete solution as the dependent variable in a univariate ANOVA with age group, distribution, and solution type as independent variables. This analysis found no main effects but did reveal a significant interaction between distribution and solution type, $F(2,173) = 12.58, p = .001, \eta^2 = .07$. For participants in the mixed distribution conditions, those with the unlikely solution needed more questions to reach the solution ($M_{\text{mixed_unlikely}} = 5.91; SD = 2.36$) than those with the likely solution ($M_{\text{mixed_likely}} = 4.16; SD = 2.10$), $t(86) = 3.673, p < .000$. However, for participants in the uniform distribution conditions, those with the unlikely solution ($M_{\text{uniform_unlikely}} = 4.23; SD = 2.03$) did not need more questions than those with the likely solution ($M_{\text{uniform_likely}} = 5.05; SD = 2.85$), $t(84) = 1.541, p = .127$. Notably, we did not find interactions between age and other variables.

Type of questions As in Study 1, questions were coded as either hypothesis scanning or constraint seeking. As in Study 1, all questions were independently coded by the experimenter and a second student assistant, both blind to the experimental hypotheses, resulting in total agreement of Kappa = .953 with $p < .001$. In the few cases where the two raters did not agree, a third Italian rater, blind to the experimental hypotheses and procedure, was consulted.

We performed a univariate ANOVA with the percentage of constraint-seeking questions asked by each participant as the dependent variable and age group, distribution, and solution type as independent variables. The analysis revealed a main effect of age group, $F(2,173) = 17.13, p < .001, \eta^2 = .18$. A Bonferroni post hoc analysis found that younger and older children asked a similar proportion of constraint-seeking questions ($M_{\text{young_child}} = 23%; SD = 37%; M_{\text{old_child}} = 20%; SD = 31%, p = 1.00$), which was lower than the proportion of constraint-seeking questions asked by young adults ($M_{\text{young_adults}} = 51%; SD = 31%, p < .001$). The analysis also revealed a main effect of distribution, $F(1,173) = 9.29, p = .003, \eta^2 = .06$. Participants assigned to the mixed distribution conditions asked a lower proportion of constraint-seeking questions ($M_{\text{mixed}} = 26%; SD = 32%$) than the participants assigned to the uniform distribution conditions ($M_{\text{uniform}} = 41%; SD = 39%$), see Figure 2. Interestingly, this effect did not interact with age ($p = .977$).

Even younger children asked a higher proportion of constraint-seeking questions when confronted with a uniform distribution ($M_{\text{uniform}} = 32\%$; $SD = 41\%$) as opposed to a mixed distribution ($M_{\text{mixed}} = 15\%$; $SD = 31\%$), $F(1,57) = 3.20$, $p = .079$, $\eta^2 = .05$. We did not find any effect of solution type.

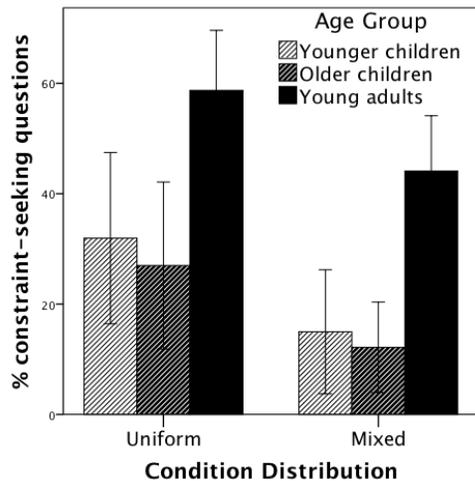


Figure 2. Study 2: Percentage of constraint-seeking questions, displayed by age group and solution type. Error bars represent one SEM in each direction.

First question We conducted a logistic regression with age group, distribution, and solution type as predictors. The Wald criterion demonstrated that both age group ($p < .001$) and distribution ($p = .011$) made a significant contribution to prediction, whereas solution type was not a significant predictor. The $\exp(B)$ value indicates that older age groups had an increased likelihood of generating an initial constraint-seeking question (by 2.94 times), after the other factors in the model were controlled for. Indeed, younger and older children asked a similar proportion of constraint-seeking questions ($M_{\text{young_child}} = 28\%$; $SD = 45\%$; $M_{\text{old_child}} = 29\%$; $SD = 46\%$, $p = 1.00$), which was lower than the proportion of constraint-seeking questions asked by young adults ($M_{\text{young_adults}} = 73\%$; $SD = 45\%$, $p < .001$). The $\exp(B)$ value also indicates that participants assigned to the mixed distribution had a decreased likelihood of generating an initial constraint-seeking question (by .42 times). Participants in the mixed distribution conditions asked a lower proportion of constraint-seeking questions ($M_{\text{mixed}} = 38\%$; $SD = 49\%$) than the participants assigned to the uniform distribution conditions ($M_{\text{uniform}} = 55\%$; $SD = 50\%$).

Discussion

Our findings replicate prior research in documenting a developmental shift from a strong tendency to ask hypothesis-scanning questions to greater reliance on constraint-seeking questions. Our findings also go beyond prior work to show, for the first time, that even young children appropriately modulate their reliance on different

question types: constraint-seeking questions become more prevalent when they yield a higher information gain and are therefore more likely to pay-off.

Studies 1 and 2 showed that both children and young adults responded to the more difficult (unlikely) version of the task by asking a higher proportion of constraint-seeking questions than they did for the easier (likely) version. While Study 1 did not find anticipated effects of scenario familiarity, in Study 2, participants asked more constraint-seeking questions when provided with alternative hypotheses that were equally likely compared to when they were given alternative hypotheses differing in probability (see Figure 2). This finding suggests that when prior knowledge (strongly) favors some hypotheses over others, participants of all ages are more likely to pursue a hypothesis-scanning strategy, possibly in the hope of achieving a quick win. However, we did find a developmental shift in participants' ability to use such prior knowledge effectively: In Study 2, young adults most often tested highly or moderately likely hypotheses first, with younger and older children significantly less likely to do so.

Study 1 confirmed the expected developmental increase in success rates. However, among participants who reached the solution, neither study found a consistent age-related boost in overall performance (assessed as the number of questions required to reach the solution). In particular, young adults' more frequent use of constraint-seeking questions did not yield a reliable advantage over younger groups, although performance was numerically (if not statistically) better for older groups when the solution was unlikely (in Study 1 and 2) or when the hypotheses provided were all equally likely (in Study 2). Adopting a constraint-seeking strategy in other conditions required giving up on the chance of obtaining a quick win by correctly guessing the solution, suggesting that a constraint-seeking strategy is not always the most efficient approach.

The rationality or optimality of a given strategy for inquiry can be defined in ecological terms (Todd et al., 2012) as the match between the structure of the task and the abilities, knowledge, experiences, and biases of the agent who has to perform the task. Our results go beyond previous demonstrations that contextual factors can influence children's strategies for inquiry by investigating the match between children's strategies and the informational structure of the problem being solved. Specifically, our design manipulated the relative *effectiveness* of different question-asking strategies, not the cognitive demands required to enact them, and our results reveal that participants did not differ in their sensitivity to these manipulations across age groups; even young children asked fewer hypothesis-scanning questions when doing so was unlikely to pay off.

What underlies children's strategies for inquiry in the real world and how do they change in response to context? One possibility is that children explicitly assess which strategy is most likely to be efficient in a given context. This seems unlikely, however, given the difficulty of making such an assessment in a real world situation, as well as the fact that

cues pointing to the likely effectiveness of a given strategy are at best indirect. We suggest instead that in real world problems, children (as well as adults) might rely on a heuristic procedure that is likely to yield flexible and efficient behavior in a range of contexts. Specifically, we propose that people ask themselves a single question: Is one of the given/self-generated alternative hypotheses sufficiently *more* likely than the others? If so, test it directly (with a hypothesis-scanning question); otherwise, collect information to reduce the number of alternative hypotheses (using a constraint-seeking question). This rough rule of thumb helps identify possible sources of developmental change: it could be that an important developmental difference lies in the initial *hypothesis-generation phase*, with young children simply generating fewer hypotheses than young adults, or in the *testing phase*, with children adopting a lower threshold for pursuing a hypothesis-scanning question. These possibilities, however, are not incompatible with other proposals: That children have not yet fully mastered the more complex constraint-seeking strategy (e.g., Denney, 1975) or that even when children know that constraint seeking would be more appropriate and efficient, they are not able to identify good constraints—that is, to abstract the object-general features that are needed to ask effective constraint-seeking questions.

More generally, the two strategies for inquiry that we consider can be seen as lying along a continuum, with hypothesis-scanning questions partitioning the search space into two very unequal partitions (the single tested solution versus *everything else*), and constraint-seeking questions achieving higher information gain the more closely they correspond to two equally-sized partitions. While the former approach has the potential for arriving at an immediate solution, the latter will tend to dominate in the long run. This trade-off might be understood in terms of “exploiting” versus “exploring” (Cohen, McClure, & Yu, 2007; Hills, Todd, & Goldstone, 2010), where hypothesis-scanning strategies *exploit* prior knowledge in the hopes of a quick win, while constraint-seeking strategies *explore* the broader hypothesis space. Understanding the current findings in these terms suggests promising directions for new research and also offers a broader framework within which the current proposal can be understood.

In sum, the process of asking questions plays a crucial role in learning and development. Nonetheless, little is known about children’s strategies for inquiry and how they change in the course of development. Here we have provided some first insights into the adaptive flexibility of children’s strategies for asking questions and proposed a tentative heuristic model. We hope that these initial steps help pave the way for additional research involving a broader range of question types and experimental paradigms and incorporating strategies for information search beyond questions, such as direct observation and experimentation.

References

- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical transactions of the Royal Society of London. Biological sciences*, *362*(1481), 933–42.
- Davidson, D. (1991a). Children’s decision making examined with an information-board procedure. *Cognitive Development*, *6*, 77–90.
- Davidson, D. (1991b). Developmental differences in children’s search of predecisional information. *Journal of Experimental Child Psychology*, *52*, 239–255.
- Denney, D. R. (1975). The effects of exemplary and cognitive models and self-rehearsal on children’s interrogative strategies. *Journal of Experimental Child Psychology*, *19*, 476–488.
- Denney, D. R., & Denney, N. W. (1973). The use of classification for problem solving: A comparison of middle and old age. *Developmental Psychology*, *9*, 275–278.
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2010). The central executive as a search process: priming exploration and exploitation across domains. *Journal of experimental psychology. General*, *139*(4), 590–609.
- Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Sciences*, *10*, 464–470.
- Mosher, F. A., & Hornsby, J. R. (1966). On asking questions. *Studies in Cognitive Growth*. New York, NY: Wiley.
- Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children’s sequential information search is sensitive to environmental probabilities. *Cognition*, *130*(1), 74–80.
- Ruggeri, A., & Feufel, M. A. (2013). *How the level of inclusiveness affects object categorization*. Manuscript submitted for publication.
- Sandberg, E. H., Huttenlocher, J., & Newcombe, N. (1996). The Development of Hierarchical Representation of Two-Dimensional Space. *Child Development*, *67*(3), 721–739.
- Siegler, R. S. (1977). The twenty questions game as a form of problem solving. *Child Development*, *48*, 395–403.
- Todd, P. M., Gigerenzer, G., & the ABC Research Group. (2012). *Ecological rationality: Intelligence in the world*. New York, NY: Oxford University Press.