

Heuristic to Bayesian: the Evolution of Reasoning from Childhood to Adulthood *

Jori Barash

Isabelle Brocas

*University of Southern California
and CEPR*

Juan D. Carrillo

*University of Southern California
and CEPR*

Niree Kodaverdian

May 2018

Abstract

In this laboratory experiment, children and teenagers learn the composition of balls in an urn through sampling with replacement. We find significant aggregate departures from optimal bayesian learning across all ages, but also important developmental trajectories. Two inference-based and two heuristic-based strategies capture the behavior of 65% to 90% of participants. Many of the youngest children (K to 2nd grade) base their decisions only on the last piece of information and use evolutionary heuristics (such as the “Win-Stay, Lose-Switch” strategy) to guide their choices. Older children and teenagers are gradually able to condition their decisions on all previous information but they often fall prey of the gambler’s fallacy. Only the oldest participants display optimal bayesian reasoning. These results are modulated by task complexity, and bayesian reasoning is evidenced earlier when inferences are simpler.

Keywords: laboratory experiment, developmental economics, learning, bayesian updating, heuristic reasoning.

JEL Classification: C91, D83, D91.

*We are grateful to members of the Los Angeles Behavioral Economics Laboratory (LABEL) for their insights and comments in the various phases of the project. We also thank the staff at the Lycée International de Los Angeles for their support, in particular Emmanuelle Acker, Nordine Bouriche, Anneli Harvey and Mathieu Mondange. All remaining errors are ours. The study was conducted with the University of Southern California IRB approval UP-12-00528. We acknowledge the financial support of the National Science Foundation grant SES-1425062. Address for correspondence: Isabelle Brocas, Department of Economics, University of Southern California, 3620 S. Vermont Ave., Los Angeles, CA 90089, USA, <brocas@usc.edu>.

1 Introduction

Everyday decisions depend on our ability to form beliefs about the surrounding environment given what we have previously observed. The ability to make correct inferences is critical to all aspects of decision-making and planning. From a descriptive point of view, the beliefs we hold at a given point in time are revised when new evidence is produced and new beliefs emerge from this process. Revision involves calculations that depend on how we interpret the evidence. Normatively, the bayesian approach provides a mathematically precise framework to determine a posterior belief from the prior belief and the evidence obtained.¹ A bayesian decision-maker holds consistent beliefs, processes all information perfectly, and always arrives at the best conclusions that can be drawn from the evidence.

Neoclassical economics has traditionally assumed that (adult) decision-makers are rational and capable of making bayesian inferences to learn about their environment. The view is supported by evidence in cognitive psychology and neuroscience which shows that the way humans assess their environment and make predictions closely follows Bayes theory. Interestingly, even neuron firing patterns are consistent with the encoding of probability distributions and posterior computations (Beck et al., 2008; Deneve, 2008; Ma et al., 2008) suggesting that bayesian mechanisms are in place even at a very basic sensory processing level. There is also evidence that children understand the notions of chance and probabilities (Denison et al., 2006; Skoumpourdi et al., 2009; Nikiforidou and Pange, 2010) and develop gradually their ability to think in probabilistic terms (Piaget and Inhelder, 1951; Kreitler and Kreitler, 1986). By the age of 4, children can combine prior knowledge about an hypothesis with new information (Gopnik and Tenenbaum, 2007; Xu and Tenenbaum, 2007; Gopnik, 2012). These findings taken together suggest that human beings are wired to process information in a bayesian way.

Yet, many studies in Economics and Psychology have reported suboptimal behavior and systematic mistakes in the context of learning in adults. Examples include the well known gambler's fallacy, a belief that, if something happens more frequently than expected during a certain period, it will happen less frequently than expected in the future. The fallacy describes how gamblers interpret a streak of failures as an indication of an increased probability of success on subsequent attempts. It emanates from the erroneous belief that small samples must be representative of the larger population and it is understood as a heuristic bias (Tversky and Kahneman, 1971, 1974; Terrell, 1994; Croson and Sundali, 2005) that generates wrong inferences and behavior inconsistent with bayesian

¹Formally, the model assumes the environment is unknown but we can form a hypothesis h about it, a prior belief $\Pr(h)$ that the hypothesis is correct, and a belief about the likelihood to observe a given piece of evidence e if the hypothesis is correct, $\Pr(e|h)$. With these elements in hand, one can formulate a posterior belief about the hypothesis h after observing evidence e as $\Pr(h|e) = \frac{\Pr(e|h)\Pr(h)}{\Pr(e)}$.

updating (Grether, 1992; Holt and Smith, 2009). Its counterpart, the hot hand fallacy, builds on the opposite and also incorrect belief that success is more likely after observing a streak of successes, and seems to play an important role in sports (Gilovich et al., 1985; Croson and Sundali, 2005).² Both fallacies refer to situations in which decision-makers take into account all the information they observe, but they do so wrongly. These have to be distinguished from cases where decision-makers use heuristic rules of behavior that do not account for proper inferences about the environment. For instance, “Win-Stay, Lose-Switch” strategies (Thorndike, 1911), which base behavior on the last revealed outcome, have been used to describe decisions in learning problems (Goodnow and Pettigrew, 1955; Steyvers et al., 2009) and sometimes have been shown to approximate behavior consistent with Bayes’ rule (Bonawitz et al., 2014). In close relation, reinforcement learning theories presuppose that a decision-maker’s goal is to select a policy that specifies which action to take after observing an outcome in order to maximize future rewards. These theories represent alternatives to models based on bayesian inferences (Erev and Roth, 1998; Camerer and Ho, 1999; Sutton and Barto, 2011) and have proved successful at describing behavior in uncertain dynamic environments. Overall, human decisions are often based on mechanistic or heuristic rules that produce decisions independent of the true information structure that underly what we observe.³

The evidence reported above underscores a tension between our wired innate ability to carry bayesian calculations and the observed tendency to make wrong inferences and/or rely on heuristic rules. This is possibly due to the differences in the paradigms employed in the literature depending on the objectives of the research. Indeed, the studies on children mostly focus on simple paradigms that help evidence sophistication in children but do not address biases in learning, whereas the literature on adults is often interested in highlighting anomalies. Therefore, it is unclear whether we are born with an ability to make correct inferences but develop wrong inference methods and mechanistic rules, or whether these anomalies are in place in children but masked by the features of the tasks selected in earlier studies.

The purpose of this paper is to study the developmental trajectory of bayesian learning and heuristic reasoning in a setting of moderate complexity that is prone to choice “mistakes.” In order to focus on the ability of our participants to revise their beliefs in a bayesian way, we consider a setting where probabilities are objective and given explic-

²Miller and Sanjurjo (2016) demonstrates, however, that there is a bias in the measure of the conditional dependence of present outcomes on streaks of past outcomes in the sequential data considered. Once the bias is corrected, the size of the fallacy is reduced.

³These decisions, however, are based on past outcomes, which are intrinsically correlated with the underlying information structure. As such, they may still generate behavior consistent with bayesian inferences (Borhani and Green, 2018).

itly rather than subjective. More precisely, we consider tasks in which an urn has one of two possible (known) compositions of green and yellow balls and we draw balls with replacement. After each draw, participants are asked to guess the color of the next drawn ball and they are rewarded if they are correct. In this framework, a bayesian learner can unambiguously compute the posterior probability of each possible composition and the likelihood that the next draw is of a given color. By contrast, a player falling prey of the gambler’s fallacy misinterprets the information revealed by a small number of draws and a heuristic user makes mechanistic guesses that neglect the informational content of all but the last draw. By asking participants to observe multiple draws and make multiple guesses, it is possible to disentangle between random play, bayesian behavior, wrong inference-based choices, and mechanistic learning. To better assess the underlying abilities of our participants, we vary task complexity by making bayesian computations more or less difficult.

We found significant aggregate departures from bayesian behavior across all ages and in both tasks. Notably, the proportion of optimal choices in school age participants was, on aggregate, at or below random level in the most complex task. Subjects also responded to complexity and we found, on aggregate, significantly higher proportions of bayesian choices in the simplest task. More importantly, a closer look at individual strategies revealed key developmental trajectories. Many of the youngest children (K to 2nd grade) used evolutionary heuristics, basing their decisions exclusively on the performance of their last choice. Older children and teenagers were gradually able to condition their choices on the whole sequence of information, but this inference-based decision making was suboptimal, in that they often fell prey of the gambler’s fallacy.⁴ Only the oldest participants displayed significant levels of optimal bayesian reasoning. Again, these results were modulated by task complexity, and bayesian reasoning was evidenced earlier when inferences were simpler. Last, strategies were correlated across tasks and the learning strategy adopted in the complex task was a predictor of the strategy in the simple one. The results indicate that behavior was driven by underlying learning capabilities. Overall, while aggregate behavior was similar in all groups, strategies evolved substantially with age, from heuristic rules, to inference-based and finally to optimal bayesian.

The paper is organized as follows. Section 2 describes the experimental design and theoretical framework. Section 3 reports our aggregate results (section 3.1), as well as the evolution of individual strategies and payoffs in the complex (section 3.2) and simple (section 3.3) tasks. It also compares the behavior of individuals across tasks (section 3.4). Section 4 gathers some concluding remarks.

⁴For other studies of heuristic reasoning in children and adolescents, see Jacobs and Potenza (1991); Baron et al. (1993); Davidson (1995).

2 Experiment

2.1 Design and procedures

Participants. We recruited 334 school-age children and adolescents –from kindergarden to 11th grade– at the Lycée International de Los Angeles (LILA), a bilingual private school in Los Angeles, with a campus in Los Feliz (pre-K to 5th grade) and another in Burbank (6th to 12th grade). We ran 35 sessions that lasted between 60 and 90 minutes.⁵ Sessions were conducted in a classroom at the school using touchscreen PC tablets and tasks were programmed in z-Tree. Sessions had 8 or 10 subjects. For each session, we mixed male and female subjects from the same grade, but for logistical reasons sometimes had to include subjects of two consecutive grades. For comparison, we ran 6 sessions with 48 USC students. These were conducted at the Los Angeles Behavioral Economics Laboratory (LABEL) in the Economics department at the University of Southern California, using identical procedures, and with subjects recruited from the LABEL subject pool.

Approximately one-half of the children in each grade participated in the experiment, except for 11th graders because a fraction of them were in a field trip. This amounts to 19 to 43 participants per grade. Given this small sample size, we decided to group subjects into four age-categories **C1** to **C4**. The partition is roughly consistent with the Piagetian stages of cognitive development emphasized in the developmental psychology literature (Piaget, 1971; Fischer and Silvern, 1985).

Finally, since this particular task has never been run with adults, we included age-category **C5** to compare the behavior of our children with that of college students. The exercise should be taken with a grain of salt due to the differences in background, ethnicity, size of the peer pool, etc. between LILA children and USC students. At the same time, it is worth noting that the majority of LILA students go to well-ranked colleges in North America and Europe, so it is not an unreasonable match.⁶ Also and unless otherwise noted, when comparing aggregate choices we perform two-sided t-tests of mean differences. We use a p-value of 0.05 as the benchmark threshold for statistical significance, and cluster standard errors at the individual level whenever appropriate. Table 1 summarizes our population and age-categories.

⁵Sessions were, on average, 10 minutes longer in elementary school because we provided extra resting time between tasks. The remaining variance in length depended on external circumstances (bringing students from the classroom, seating them, choosing the rewards, etc.). Overall, sessions were slightly longer than optimal and fatigue may have played a role for some of the younger children. We however believe that the effect was small: fatigue should have pushed students to play close to random or to use patterns, a behavior that we did not observe in the data.

⁶The majority of studies with children do not perform the same experiment with an adult population. We believe it is valuable to include an adult control group to establish a baseline, while acknowledging the important limitations due to potential differences in personal characteristics.

Location	LILA Los Feliz		LILA Burbank		USC
Age-category	C1	C2	C3	C4	C5
Grade	K-1 st -2 nd	3 rd -4 th -5 th	6 th -7 th -8 th	9 th -10 th -11 th	
# subjects	30-19-30	29-29-19	43-40-31	31-21-12	48

Table 1: Subjects by grade

Tasks. The experiment consists of three sets of tasks always performed in the same order. The first set of tasks is a series of one-shot binary-choice anonymous dictator games. The second set of tasks consists of two binary-choice dynamic, alternating dictator supergames, with a fixed and anonymous partner. The objective of these tasks is to study altruism and strategic giving. Since they are designed to study a different paradigm, we relegate the analysis to a different paper (Brocas et al., 2017).

The third set of tasks –the ones studied in this paper– is a learning exercise. In these tasks, balls are sequentially drawn with replacement from an urn with an uncertain number of yellow (Y) and green (G) balls. Participants are asked to predict the color of each upcoming ball and are rewarded for correct predictions.

To be more precise, we designed a simple but engaging protocol, which is identical in all sessions and age-groups. In task 1 participants are presented with a physical box containing six yellow balls and six green balls (6Y&6G). The experimenter removes one ball without looking at it or showing it to anyone and puts it in his pocket. Participants are explicitly instructed that the eleven remaining balls are equally likely 6Y&5G or 5Y&6G. Three participants take turns to draw one ball each without looking in the box, show the ball to the room and put it back in the box. After the three draws, participants are asked to predict the color of the next ball (yellow or green), using a touchstone tablet as interface. Then, a participant draws one ball without looking, shows it to the room and puts it back. The experiment records in the computer server the outcome of each draw, and each client computer provides feedback to the participant on whether his/her prediction was correct. This procedure is repeated a total of 6 times. A graphical representation of task 1 is provided in Fig.1 (left).⁷

Task 2 follows the exact same procedures except that prior to removing the first ball, the experimenter removes one other ball, shows it to the room and leaves it on the table. If the ball on the table is green, task 2 is identical to task 1 except that at the time where participants draw balls, the urn contains either 6Y&4G or 5Y&5G.⁸ A graphical

⁷We use a physical environment and ask participants to help drawing balls in order to increase both their engagement and their trust on the procedure.

⁸In other words, participants first draw three balls. For the following six rounds, they make a prediction

representation is provided in Fig.1 (right). Naturally, if the ball on the table is yellow, colors are flipped and task 2 is identical to task 1 except that at the time where participants draw balls, the urn contains either 5Y&5G or 4Y&6G. Task 2 always follows task 1.




Figure 1: Learning task. Task 1 (left) and Task 2 (right).

Notice that we decided to include in the box 6 balls of each color (instead of just 2 or 3 for example) because, otherwise, participants with suboptimal predictions might accumulate very few tokens and get frustrated. In other words, the higher the number of initial balls, the higher the impact of chance and the lower the impact of optimal decision making. We felt this was important in an experiment with children but would have not been our choice had we run the experiment only with adults.

Payoffs. During the experiment, subjects accumulate tokens in all tasks and rounds.⁹ We implemented two different conversions depending on the subjects' ages. USC students and subjects at LILA Burbank (grades 6th to 11th) had tokens converted into money, paid with an Amazon gift card at the end of the experiment.¹⁰ For subjects at LILA Los Feliz (grades K to 5th) we set up a shop with 20 to 30 pre-screened, age appropriate toys.¹¹ Different toys had different token prices. Before the experiment, children were taken to the shop and showed the toys they were playing for. They were also instructed about the token prices of each toy and, for the youngest subjects, the experimenter explicitly stated that more tokens would result in more toys. At the end of the experiment, subjects and draw a ball.

⁹We paid subjects for all the tokens accumulated. For a discussion of the relative merits of paying subjects for all, some or one round, see Charness et al. (2016).

¹⁰The conversion rate for USC subjects (\$0.15/token) is higher than for LILA Burbank subjects (\$0.07/token) to correct for differences in marginal value of money and opportunity cost of time. It implied large differences in average earnings (\$22.3 vs. \$10.0) despite similar average number of tokens obtained (149 vs. 143). In compliance with LABEL policies, USC subjects were also paid a \$5 show-up fee.

¹¹These included gel pens, friendship bracelets and erasers for young girls, figurines, die-cast cars and trading cards for young boys, and apps, calculators and earbuds for older kids. However, children were free to choose any item they liked within their budget.

learned their token earnings and were accompanied to the shop to exchange tokens for toys. We made sure that every child earned enough tokens to obtain at least three toys. At the same time, no child had excess tokens after choosing all the toys they liked.¹² A transcript of the instructions read aloud is included in Appendix A.¹³

2.2 Theory

Optimal decision. Suppose there are two urns, A and B , each with likelihood $p(A)$ and $p(B)$ ($= 1 - p(A)$). Urn A contains a_Y yellow balls and a_G green balls. Urn B contains b_Y yellow balls and b_G green balls. The probability of drawing a yellow ball from urns A and B are $q_Y^A = \frac{a_Y}{a_Y + a_G}$ and $q_Y^B = \frac{b_Y}{b_Y + b_G}$. The probability of drawing a green ball from urns A and B are $q_G^A = 1 - q_Y^A$ and $q_G^B = 1 - q_Y^B$.

Suppose balls are drawn with replacement and we observe n_Y yellow and n_G green balls. Using Bayes rule we can compute the posterior probability that the urn is A as:

$$\begin{aligned} p(A | n_Y, n_G) &= \frac{p(n_Y, n_G | A) p(A)}{p(n_Y, n_G | A) p(A) + p(n_Y, n_G | B) p(B)} \\ &= \frac{1}{1 + \left(\frac{q_Y^B}{q_Y^A}\right)^{n_Y} \left(\frac{q_G^B}{q_G^A}\right)^{n_G} \frac{p(B)}{p(A)}} \end{aligned}$$

Given (n_Y, n_G) , it is optimal to predict that the next draw is yellow if and only if:

$$U_Y(n_Y, n_G) \equiv p(A | n_Y, n_G) q_Y^A + p(B | n_Y, n_G) q_Y^B > \frac{1}{2} \quad (1)$$

Let urns A and B correspond to the case where the ball removed (and not shown) is yellow and green respectively. By construction, in our experiment $p(A) = p(B) = 1/2$.

In task 1, this means that $a_Y = b_G = 6$ and $a_G = b_Y = 5$, so $q_Y^A = q_G^B = 6/11$ and $q_G^A = q_Y^B = 5/11$. After some algebra, condition (1) becomes:

$$\left(\frac{5}{6}\right)^{n_Y - n_G} < 1 \quad \Leftrightarrow \quad n_Y > n_G$$

For task 2, suppose that the ball initially removed and shown to everyone is green (the case of yellow is symmetric). After removing (and not showing) the second ball, we have

¹²The procedure emphasized the value of earning tokens but, at the same time, ensured an enjoyable and exciting experience.

¹³Notice that our instructions do not contain a comprehension quiz to avoid a stressful, test-like situation. Instead, we orally asked questions to make sure students understood, while remaining casual and interactive.

$a_Y = 6$, $a_G = 4$, $b_Y = 5$ and $b_G = 5$, so $q_Y^A = 3/5$, $q_G^A = 2/5$ and $q_Y^B = q_G^B = 1/2$. Condition (1) becomes:

$$q_Y^A > \frac{1}{2}$$

which is, by construction, true for all n_Y and n_G .

Finally, notice that if $q_Y^A > 1/2 > q_Y^B$, then $p(A | n_Y, n_G)$ is increasing in n_Y and decreasing in n_G , and so is $U_Y(n_Y, n_G)$. The optimal decision is summarized in Table 2.

Task	Task 1	Task 2	
	6Y&5G or 5Y&6G	6Y&4G or 5Y&5G	5Y&5G or 4Y&6G
Decision	Y if $n_Y > n_G$ G if $n_Y < n_G$	Y for all n_Y, n_G	G for all n_Y, n_G

Table 2: Optimal decision

Interpretation. Although the mathematical derivation is not trivial, the intuition is straightforward. In task 1, the urn is equally likely to contain 6Y&5G or 5Y&6G. If the difference between the observed number of yellow and green draws is positive (negative), then the urn is more (less) likely to contain 6Y&5G than 5Y&6G, implying that the next draw is more likely to be yellow (green). As the absolute difference between yellow and green draws increases, the confidence in the color of the next draw also grows.

The reasoning is different for task 2. Suppose a green ball is on the table, that is, the urn initially contains 6Y&5G. After removing a ball, the urn has either more yellow than green balls (6Y&4G) or an equal number of yellow and green balls (5Y&5G). A yellow draw is therefore (weakly) more likely than a green draw independently of the observed past sequence (the reasoning is obviously symmetric for 5Y&6G).

Given the mathematical complexity to find the optimal decision rule using bayesian probability theory, we do not expect participants of any age to perform such calculation. At the same time, we selected this design because the optimal strategy is rather simple. We therefore expect that some participants will play the optimal (or close to optimal) strategy based on a non-mathematical, non-analytical, intuitive understanding of the situation.

3 Results

3.1 Children do not play (close to) optimally

As noted above, Bayes theory states that a rational participant should make a prediction as a function of the difference between yellow and green past draws in task 1, and independently of past draws in task 2. We compute for each participant in each age-category

described in Table 1 the proportion of choices that match and do not match the prediction of bayesian theory, independently of whether the choices turned out to be correct or not.¹⁴ Fig.2 presents different cuts of this data.

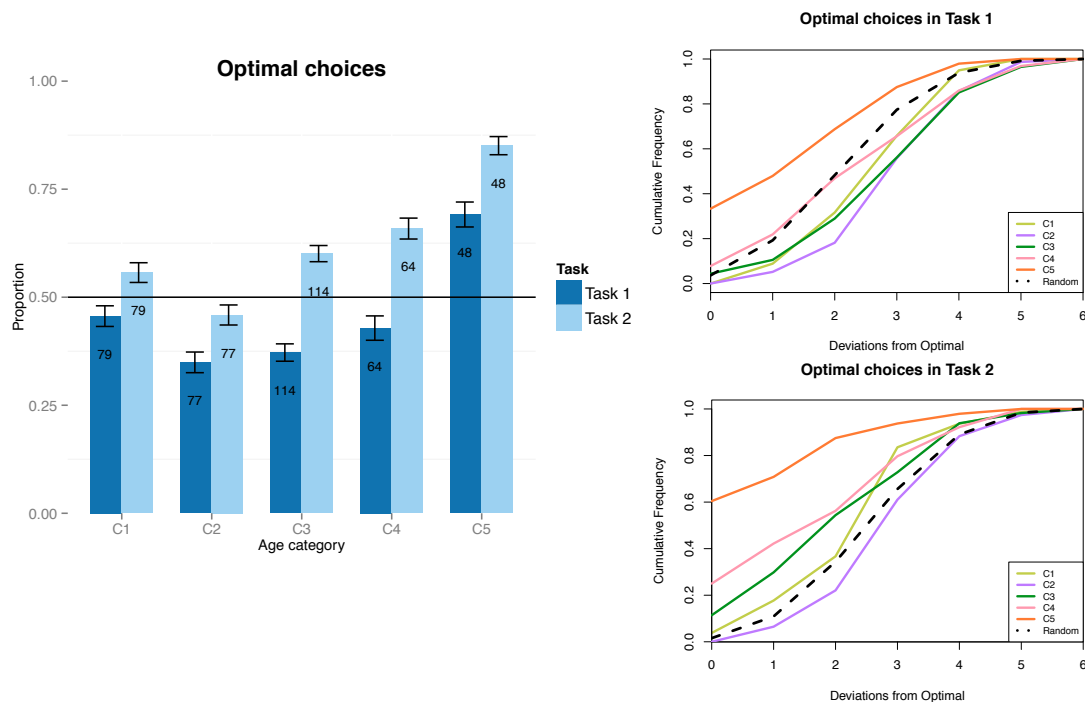


Figure 2: Optimal behavior. Average proportion of bayesian choices in tasks 1 and 2 by age-category [left] (random = 0.5). Cumulative distribution function of the number of deviations from bayesian choices by age-category in task 1 [right-up] and task 2 [right-down] (c.d.f. under random choice in dashed black).

As we can see from Fig.2 (left), if we look at the behavior of school-age participants (**C1** to **C4**), we notice that the average proportion of optimal choices in task 1 is statistically below what we would observe under random behavior (0.5) in all age-categories except **C4** (t-test, p-value < 0.05). Optimal choices are also statistically lower in **C2** and **C3** than in **C1** and in **C2** than in **C4** (t-test, p-value < 0.05). By contrast, our college students **C5** perform better than random and better than any other age-category (t-test, p-value < 0.001).

The proportion of optimal choices is higher in task 2 than in task 1 in all school-age categories as well as in the adult population (t-test, p-value < 0.01). All school-age

¹⁴For task 1, we do not include observations where the number of past yellow and green draws is equal, since either choice is optimal.

categories except **C2** outperform random choice in task 2 (t-test, p-value < 0.01) and the pattern across age is the same as for task 1: average performance is U-shaped and statistically lower in **C2** than in **C1**, **C3** and **C4** (t-test, p-value < 0.01). Finally, only college students outperform all other ages and exhibit the performance expected from a population with a majority of bayesian players (85.1% of optimal choices). The results hold if we correct for multiple comparisons using the Holm method. In Appendix B, we present the same information, except that we group children by grade. The number of observations is reduced, thereby increasing the confidence intervals around the mean. However, results are qualitatively very similar, with the exception of 6th and 11th graders who perform relatively better in tasks 2 than the children in their neighboring grades.

These results are confirmed when we look at the entire distribution of equilibrium actions. For task 1 (Fig.2, right-up) the c.d.f. of the number of deviations from optimal choice stochastically dominates random behavior for subjects in **C5** (Kolmogorov-Smirnoff test, p-value < 0.001 and Mann-Whitney test, p-value < 0.001) and it is stochastically dominated by random behavior for subjects in **C1**, **C2**, and **C3** (Kolmogorov-Smirnoff test, p-value < 0.05 and Mann-Whitney test, p-value < 0.001). For task 2 (Fig.2, right-down) the c.d.f. of all age-categories except **C2** dominates random behavior (Kolmogorov-Smirnoff test, p-value < 0.01 and Mann-Whitney test, p-value < 0.01) whereas the c.d.f. of **C2** is dominated by random behavior (although only under the Mann-Whitney test, p-value < 0.05). Again, only **C5** in task 2 exhibit a behavior close to optimal (60.4% of participants with 0 deviations and 87.5% of participants with 2 deviations or less).

A possible explanation for the relatively poor performance in task 1 is that when the difference between the number of observed green and yellow past draws is small, the next draw is “almost equally” likely to be of either color. In other words, the confidence in the color of the upcoming draw is proportional to the absolute difference in the number of past green and yellow draws (see section 2.2. for the formal argument). Suppose now that subjects play optimally but make some “mistakes”, for example because attention is costly. If the cost of attention is constant, mistakes should be more frequent when the marginal benefit of taking the best action is smaller. This would mean that we should observe a higher proportion of optimal choices when one color has been drawn substantially more often than the other. To test this theory, we report in Fig.3 the proportion of optimal choices in task 1 as a function of the absolute difference in the number of past green and yellow draws.¹⁵ Due to the limited number of observations, we group **C1** with **C2** and **C3** with **C4**. We do not report decisions when the difference is 0 because, as already noted,

¹⁵In other words, we combine the probability that the participant predicts yellow when he has observed x more yellow than green balls with the probability that he predicts green when he has observed x more green than yellow balls.

both choices are optimal. Finally and again due to the small number of observations, we group in the category 4+ the cases where the absolute difference between yellow and green draws is 4 or higher.

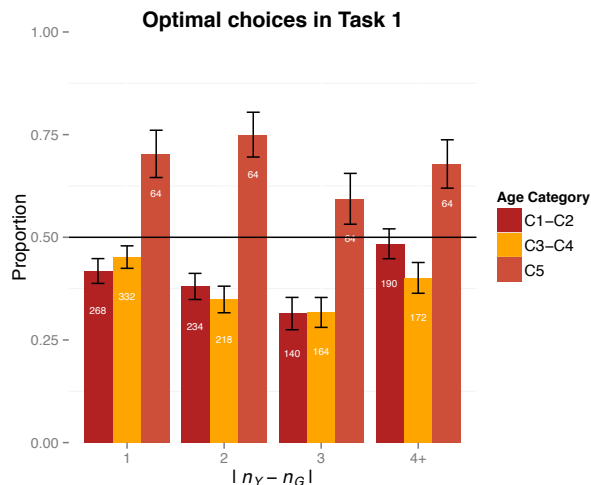


Figure 3: Choices given past draws in task 1. Average proportion of optimal choices in task 1 by age-category as a function of the absolute difference between the number of yellow and green draws in the past $|n_Y - n_G|$.

As we see from Fig.3, there is no significant trend in the proportion of optimal behavior as a function of the absolute difference between the number of yellow and green past draws. School-age participants perform below chance in six out of eight cases (Binomial test, p -value < 0.05). By contrast, college students perform significantly above chance when $|n_Y - n_G|$ is 1 or 2 (Binomial test, p -value < 0.05) but not when it is 3 or 4+. In other words, subjects do not choose more often the action predicted by bayesian theory if the evidence in favor of that action is stronger. Instead, it appears that they either realize and apply the optimal rule or they consistently follow a different strategy.

To further investigate optimal choices, we ran two Ordinary Least Squares (OLS) regressions of the proportion of correct choices of each participant in each task on age dummies for each age category **C2** to **C5** and gender as well as an OLS regression on the full sample including a task dummy ($Task2 = 1$ for observations from task 2) as regressor. The results are reported in Table 3.

This set of regressions confirms our previous results on the U-shaped average performance with age. Compared to **C1**, the proportion of optimal choices in both tasks is lower in **C2** and higher in **C5**. For task 1, optimal choices are also lower in **C3** than in **C1**, whereas for task 2 optimal choices are also higher in **C4** than in **C1**. Also, according to

	<i>Task 1</i>		<i>Task 2</i>		<i>All tasks</i>
	(1)	(2)	(3)	(4)	(5)
C2	-0.116** (0.038)	-0.119** (0.038)	-0.098** (0.036)	-0.101** (0.037)	-0.110** (0.023)
C3	-0.088* (0.035)	-0.088* (0.035)	0.044 (0.033)	0.044 (0.033)	-0.022 (0.024)
C4	-0.030 (0.040)	-0.030 (0.040)	0.102** (0.038)	0.102** (0.038)	0.037 (0.032)
C5	0.224** (0.044)	0.226** (0.044)	0.294** (0.042)	0.295** (0.042)	0.260** (0.034)
<i>female</i>	—	-0.020 (0.025)	—	-0.016 (0.024)	-0.018 (0.019)
<i>Task2</i>	—	—	—	—	0.172** (0.016)
<i>constant</i>	0.459** (0.027)	0.470** (0.030)	0.557** (0.026)	0.566** (0.029)	0.432** (0.020)
Observations	382	382	382	382	764
R-squared	0.161	0.163	0.201	0.202	0.265

* $p < 0.05$; ** $p < 0.01$; (st. errors in parenthesis)

Table 3: OLS regressions of optimal choices in tasks 1 and 2

regression (5) the proportion of optimal choices is associated with task complexity.

3.2 Task 1: from evolutionary heuristics to (some) bayesian thinking

While the above result establishes significant aggregate departures from optimal behavior in task 1, it does not describe which other strategies the participants may be following. From the previous analysis, it is also not possible to determine the degree of heterogeneity in behavior across individuals. To understand individual behavior, we retain six possible strategies commonly discussed in the literature (Croson and Sundali, 2005; Steyvers et al., 2009; Bonawitz et al., 2014), that pertain to three different modes of reasoning and cover a reasonably comprehensive spectrum of options.

Inference-based strategies. These are strategies that incorporate all the information. In particular, current decisions depend on the full string of past draws and, only indirectly, on past successes or failures. We consider two such strategies. (1) *Optimal1* is the bayesian strategy that a payoff-maximizing participant should follow in task 1. It

consists of choosing the color that has been drawn most often in the past, independently of the outcome. (2) *Reverse* is the strategy of a subject who draws inferences from all information but falls prey of the gambler’s fallacy. It consists of choosing the color that has been drawn least often in the past. By definition, *Optimal1* and *Reverse* yield highest and lowest expected payoffs. Naturally, due to the small number of observations and the inherent stochasticity of outcomes, chance is likely to play an important role in actual payoffs.

Evolutionary heuristics. In these strategies, current decisions depend exclusively on the success or failure of the last choice. We consider two cases. (3) *WStay*, or Win-Stay Lose-Shift, consists of repeating color after a success and changing after a failure. (4) *WShift*, or Win-Shift Lose-Stay, is the opposite strategy that consists of changing color after a success and repeating after a failure.¹⁶ These strategies are cognitively less demanding than inference-based strategies. They are based on outcomes (and therefore only indirectly on information) and markovian, as they do not require to track down the entire sequence. This is the main reason why we label them “heuristics.”¹⁷ Finally, notice that payoffs are likely to be higher with *WStay* than with *WShift*. Since winning is an imperfect indication that the color chosen is more prevalent than the other, it is statistically better to repeat a color after a success than to switch.¹⁸

Simplistic choices. Current decisions are based on visually appealing patterns or sequences. We consider two such strategies. (5) *Color* consists of choosing always the same color. (6) *Alternate* consists of alternating between the two colors. These strategies are eye-catching and visually pleasant but lack any cognitive or optimization basis.

For each subject, we look at the sequence of the nine draws and six choices, and determine how many of the choices coincide with the prediction of each strategy.

The left graph of Fig.4 presents a Venn diagram that depicts the number of subjects who play according to each strategy. To accommodate small “mistakes”, we include subjects who perfectly conform to a strategy and those who deviate once.^{19,20} Notice that

¹⁶ *WStay* coincides with a strategy where the subjects chooses the same color as the last ball drawn and *WShift* coincides with a strategy where the subjects chooses the opposite color of the last ball drawn. Our design cannot disentangle between these different heuristic rules.

¹⁷The reader might have in mind a slightly different definition and also view *Reverse* as a heuristic. Indeed, it is consistent with the representativeness heuristic (Kahneman and Tversky, 1972). However, we are mostly interested in distinguishing between strategies that require inferences from information (such as *Reverse*) and strategies that do not.

¹⁸We do not consider more sophisticated reinforcement learning strategies because those models require a large number of observations to be fit.

¹⁹Remember that *Optimal1* and *Reverse* make no prediction when the number of past yellow and green draws are equal. We do not count those observations as consistent or not consistent with the strategy.

²⁰There are two options for *Color* (always green or always yellow) and two options for *Alternate* (starting

some sequences of choices are compatible with more than one strategy (although a sequence can never be simultaneously consistent with (1) and (2), with (3) and (4) or with (5) and (6)). The Venn diagram accounts for this possibility by locating the individual at the intersection of all the strategies compatible with the choices.²¹ No subject is uniquely classified as *Color* and 5 subjects are uniquely classified as *Alternate*. This corresponds to 0% and 1.3% of the sample, whereas the likelihood of playing these strategies by chance is $2/2^6 \simeq 3\%$. For parsimony, we do not include these strategies in the diagram.

The right graph of Fig.4 summarizes the percentage of subjects within each age-category who conforms to the different strategies, ordered from highest to lowest expected number of optimal choices (*Optimal1* to *Reverse*, bottom to top). To facilitate comparisons, we group together all the strategies consistent with *Optimal1* and call them *O1* (*Optimal1*; *Optimal1/WStay*; *Optimal1/WShift*). We also group together all the strategies consistent with *Reverse* and call them *R* (*Reverse*; *Reverse/WShift*).²² The statistical tests reported below are Pearson’s chi-square tests of comparison of proportions.

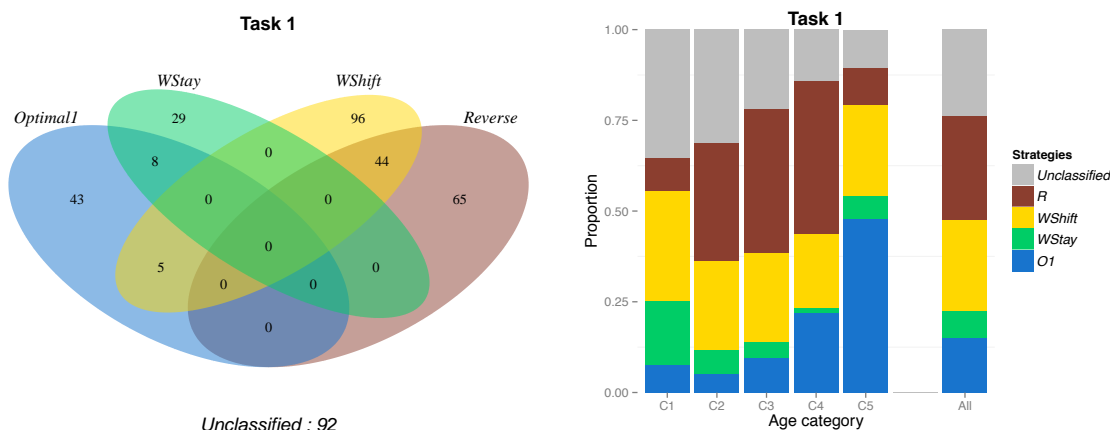


Figure 4: Task 1. Left: Venn diagram with number of participants who play according to each strategy (individuals at the intersection of two strategies have a behavior consistent with both). Right: proportion of participants within each age-category who conforms to these strategies, ordered from lowest to highest expected number of optimal choices.

with green or starting with yellow). For comparability with the other strategies, in the case of these two strategies we do not allow any deviation.

²¹Subjects are lexicographically classified based on the closest strategy. So, if a behavior is compatible with one strategy given no deviation and another strategy given one deviation, then it classifies the individual only in the strategy given no deviation.

²²An implication of this grouping method is that inference-based strategies take precedence over heuristic-based strategies. Notice however that the only intersection of strategies with a significant number of subjects is *Reverse/WShift*.

The first thing to notice is that *Optimal1*, *WStay*, *WShift* and *Reverse* account for a large fraction of decisions: the choices of 65% to 90% of individuals can be explained by one or a combination of these four strategies, with fewer *Unclassified* subjects in the older age-categories. This, however, does not mean that participants maximize payoffs or play similarly across ages. From Fig.4 (right), we notice that the youngest children (**C1**) play mostly the evolutionary heuristics *WStay* and *WShift*. We then notice a significant increase in *R* ($p < 0.001$), a strategy that remains equally popular from **C2** to **C4** and a significant decrease in *WStay* and *WShift* ($p = 0.045$ between **C1** and **C2**). We observe also a significant and gradual decrease of *Unclassified* starting in **C3**. From **C3** to **C5** there is a significant increase in *O1* ($p = 0.043$ between **C3** and **C4**, $p = 0.007$ between **C4** and **C5**).²³ Finally, the college population **C5** exhibits the largest levels of *O1* and lowest levels of *R* across all groups, consistent with their substantially higher proportion of optimal choices than any other age-category. Overall, the results emphasize the importance of the individual-level analysis. Indeed, while all school-age subjects are similar in their (below chance) proportion of optimal choices (Fig.2), the composition of the strategies played by the subjects changes significantly with age.

Fig.5 reports the average proportion of optimal choices as a function of the strategy employed by the subject, and grouping all age-categories together.

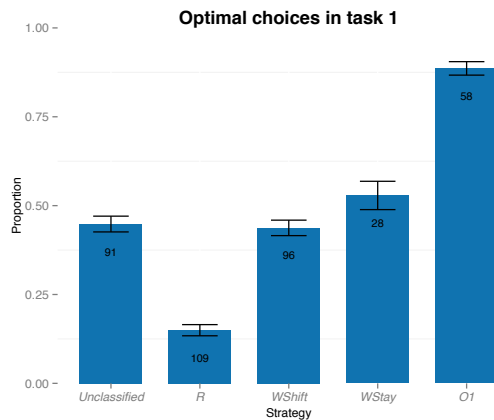


Figure 5: Task 1. Average proportion of optimal choices by strategy.

By definition, strategies consistent with *Optimal1* have all or all but one optimal choices, whereas strategies consistent with *Reverse* have none or one optimal choices. Consistent with the idea that success is more likely when the subject takes the correct action than when he does not, there are more optimal choices under *WStay* than under

²³If we correct for multiple comparisons, the trend is the same but less sharp.

WShift. Although the difference is statistically significant ($p < 0.01$), it is small in magnitude (0.53 vs. 0.44), suggesting that outcomes are not very likely to induce learning from one evolutionary heuristic to the other. Rates of optimal choices for the *Unclassified* strategies are also close to random.

All in all, the combination of Figs. 4 and 5 helps explaining the aggregate findings presented in Fig.2. Our youngest subjects **C1** use mostly evolutionary heuristics (where behavior is shaped by success or failure of the last decision) or strategies that are hard to classify (but not pattern based), resulting in average levels of optimal choices that are close to random. **C2** and **C3** start conditioning their behavior on the entire sequence of draws (as bayesian thinkers would do) except that they fall prey of the gambler’s fallacy, so they end up decreasing rather than increasing the number of optimal choices. Many **C4** keep falling on the gambler’s fallacy but some of them start reasoning like optimal bayesian players. Finally, a majority of **C5** understand and play the bayesian strategy, with the corresponding substantial increase in optimal decision making. Perhaps surprisingly, in all age-categories there is a stable fraction of subjects (around 25%) who play the counterintuitive strategy *WShift*. It is important to highlight that despite the close to random average behavior in all school-age subjects, the vast majority of individuals follow some well-defined and reasonable (albeit suboptimal or even payoff-minimizing) choice rules. Heterogeneity in strategies is responsible for the small differences in aggregate rates of optimal decision-making across individuals.

Finally, we perform a set of regressions to study the determinants of equilibrium and non-equilibrium behavior in our population. More precisely, we run Probit regressions on the strategies of the participants, where the dependent variable is whether the subject follows strategy *O1* (regressions (1) and (2) - *O1*) and whether the subject follows strategy *R* (regressions (3) and (4) - *R*). For the regressors, we include dummies for each age category **C2** to **C5**, gender and *3same* a dummy that takes value 1 if the first three draws in the session (before participants need to provide their first prediction) are of the same color. The idea behind this last variable is that three identical draws facilitate the selection of the optimal rule *Optimal1*, which is based on $|n_Y - n_G|$. The results are summarized in Table 4.

Confirming previous results, there is an evolution with age from heuristic choices to inference-based behavior (both the payoff-maximizing *O1* and the payoff-minimizing *R*). Indeed, **C4** and **C5** play the optimal strategy significantly more often than the youngest children (regressions (1) and (2)), while **C2**, **C3** and **C4** play the payoff-minimizing strategy more often than the youngest children (regressions (3) and (4)). Gender and repeated draws of the same color do not seem to have a strong and systematic effect on behavior.

	<i>O1</i>		<i>R</i>	
	(1)	(2)	(3)	(4)
C2	-0.193 (0.316)	-0.212 (0.322)	0.895** (0.248)	0.848** (0.250)
C3	0.131 (0.264)	0.159 (0.270)	1.082** (0.232)	0.997** (0.235)
C4	0.656* (0.272)	0.690* (0.275)	1.152** (0.254)	1.102** (0.255)
C5	1.381** (0.276)	1.356** (0.289)	0.091 (0.315)	0.186 (0.326)
<i>female</i>	—	-0.245 (0.171)	—	0.118 (0.145)
<i>3same</i>	—	0.133 (0.196)	—	-0.338* (0.172)
<i>constant</i>	-1.433** (0.209)	-1.366** (0.243)	-1.349** (0.199)	1.273** (0.223)
Observations	382	382	382	382
AIC	290	291.4	426.5	425.9

* $p < 0.05$; ** $p < 0.01$; (st. errors in parenthesis)

Table 4: Probit regressions of strategies in task 1 by age.

3.3 Task 2: a sharper evolution of behavior across ages

We conduct the same analysis of individual strategies in task 2. *Optimal2*, the payoff maximizing strategy in task 2, consists now of choosing always the same color (yellow if the urn starts with 6Y&5G and green if the urn starts with 5Y&6G).²⁴ *Reverse*, *WStay*, *WShift* and *Alternate* are defined exactly as before. Notice that *Reverse* is not the opposite of *Optimal2*, therefore it does not minimize payoffs. Indeed, the payoff-minimizing strategy is to choose always the same color, and opposite to the color in *Optimal2*. Only 1 subject is uniquely classified in this strategy. Also, only 3 subjects are uniquely classified as

²⁴Interestingly, from an empirical standpoint *Optimal2* often coincides with *Optimal1*, the optimal strategy in task 1 (choose the color that has been drawn most often in the past). Indeed, with more balls of one color in the urn (say, yellow), it is statistically likely to observe a sequence where yellow draws always outnumber green draws, in which case both strategies prescribe the same behavior (choose always yellow). This similitude is reflected in our sample, where 66% of subjects classified as *Optimal2* have a behavior also consistent with *Optimal1*. It means that some of the subjects that we label as bayesian optimizers in task 2 may in fact be following the bayesian optimizing strategy of the previous task. While the distinction is subtle and worth noting, we believe it does not contradict the main message of the paper.

Alternate. We again drop these two rarely played strategies from the analysis.

Fig.6 is the analogue of Fig.4 for task 2. For the right graph, we again group together all the strategies consistent with *Optimal2* and all the strategies consistent with *Reverse*, and call them *O2* and *R* respectively.²⁵

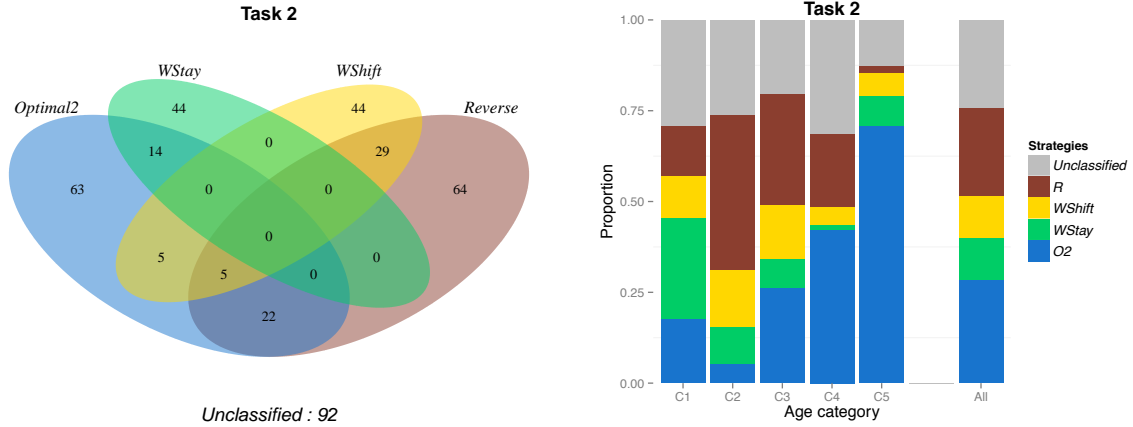


Figure 6: Task 2. Left: Venn diagram (all age-categories together). Right: distribution of strategies across age-categories.

Just like in task 1, the strategies considered in task 2 account for the majority of choices of our subjects (71% to 88% of the sample depending on the age-category). We notice a similar though sharper evolution in the behavior across age-categories. As before, many **C1** subjects (39%) play the evolutionary heuristics *WStay* and *WShift*. **C2** make significantly fewer optimal choices compared to **C1**; they also increase *R* and decrease *WStay* ($p < 0.05$). We then observe a substantial and sustained increase in *O2* ($p < 0.001$ between **C2** and **C3**, $p = 0.036$ between **C3** and **C4**, $p = 0.002$ between **C4** and **C5**), at the expense of *WShift* and *WStay* ($p = 0.02$ between **C1** and **C3**, $p = 0.004$ between **C2** and **C4** and $p = 0.009$ between **C3** and **C4**) then at the expense of *R* ($p = 0.009$ between **C4** and **C5**) and *Unclassified* ($p = 0.036$ between **C4** and **C5**). The level of optimal play reaches a maximum of 71% in the adult population.

It is also instructive to compare the evolution of strategies across tasks. There is an increase between *O1* and *O2* in categories **C3**, **C4** and **C5**. This change is significant (differences of 16.7%, 20.3% and 22.9%, $p < 0.04$). These increases come at the expense of decreases in *WShift* in **C4** and **C5** and also at the expense of a decrease in *R* in **C4**. There are at least three possible reasons for this increase in bayesian choices. First,

²⁵To give bayesian decision making a chance, subjects at the intersection of *Optimal2* and *Reverse* are put in *O2*. Classifying them as *R* does not alter significantly the results.

optimal behavior is arguably easier in task 2 (choose always the same color). It also coincides with the computation of a belief with no subsequent revision. Second, in task 2 optimal behavior and the behavior which was optimal in task 1 but is not anymore often coincide (see footnote 24), providing additional chances to observe such a choice. Finally, tasks are always performed in the same order, and some participants may have transferred knowledge on optimal behavior from task 1 to task 2. Unfortunately, the experiment is not designed to disentangle between these hypotheses.

We also report in Fig.7 the proportion of optimal choices for each strategy using the same methodology as in task 1.




Figure 7: Task 2. Average proportion of optimal choices by strategy.

We obtain the same ordering as before. By construction, strategies consistent with *Optimal2* have all or all but one optimal choices. Choices are again close to random (and, this time, not significantly different from each other) under *WStay* and *WShift* (0.52 and 0.50). Since *Reverse* is not payoff-minimizing, the proportion of optimal choices under *R* is substantially higher than in task 1 (0.43). It is statistically lower than *WStay* (p-value < 0.05) but not statistically lower than *WShift*.

Overall, the evolution across age-categories observed in task 2 is qualitatively similar but quantitatively sharper than that observed in task 1. Evolutionary heuristics start high but decrease with age. Children in **C2** are prey of the gambler’s fallacy. As they grow, they slowly understand what the optimal strategy is, and by the adult age they consistently play according to it.

We also performed the same set of Probit regressions as in task 1. The results are collected in Table 5.

The evolution of behavior with age in task 2 –from heuristic to bayesian– is similar than in task 1: as before, **C4** and **C5** exhibit a significant increase in the optimal strategy

	<i>O2</i>		<i>R</i>	
	(1)	(2)	(3)	(4)
C2	-0.700* (0.290)	-0.693* (0.291)	0.904** (0.227)	0.902** (0.236)
C3	0.292 (0.208)	0.326 (0.211)	0.579** (0.214)	0.507* (0.221)
C4	0.729** (0.229)	0.795** (0.235)	0.253 (0.250)	0.088 (0.259)
C5	1.475** (0.253)	1.580** (0.266)	-0.953* (0.447)	-1.223** (0.461)
<i>female</i>	—	-0.041 (0.149)	—	-0.286 (0.154)
<i>3same</i>	—	0.236 (0.179)	—	-0.911** (0.198)
<i>constant</i>	-0.926** (0.165)	-1.006** (0.202)	-1.084** (0.176)	-0.673** (0.202)
Observations	382	382	382	382
AIC	391.8	394	393.9	371.4

* $p < 0.05$; ** $p < 0.01$; (st. errors in parenthesis)

Table 5: Probit regressions of strategies in task 2 by age.

whereas **C2** and **C3** show a significant increase in *Reverse* at the expense of heuristic reasoning. However and as already noted, the evolution is more pronounced than before, as we also observe a statistically significant decrease in *O2* for **C2** as well as a significant decrease in *R* for **C5**. We also notice fewer suboptimal strategies when the first three draws are of the same color.

3.4 From task 1 to task 2: sticking to a strategy

We next investigate the relationship between behavior in task 1 and task 2. Figure 8 provides heatmaps with the strategies followed by individuals in both tasks, where row i corresponds to strategy i in task 1 and column j to strategy j in task 2. In the left graph, each cell represents p_{ij} , the proportion of subjects who follow the pair of strategies (i, j) .

In the right graph, each cell represents $p_{ij} / \sum_j p_{ij}$, the proportion of subjects who follow strategy j in task 2 among all the subjects who follow strategy i in task 1. In both cases, cells along the 45-degree line represent different measures of individuals who play the same strategy in both tasks. Because of the limited data, for the analysis in this section we group together subjects of all age-categories.

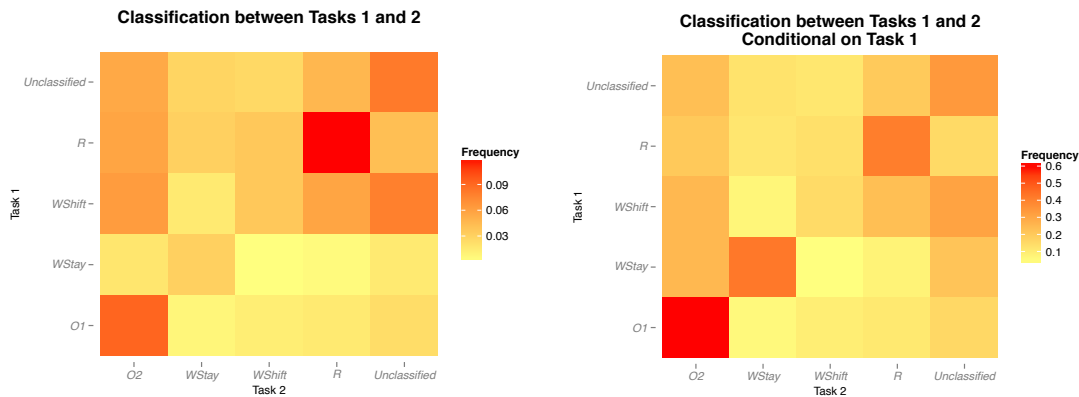


Figure 8: Strategies in tasks 1 and 2: unconditional (left) and conditional on behavior in task 1 (right).

From the left graph, we notice a significant correlation of behavior across tasks among subjects who play the inference-based strategies, *Optimal* or *Reverse*. Subjects who use a non-discernible strategy in one task are also more likely to remain unclassified in the other. The correlation is less pronounced for the evolutionary heuristic *WStay* and non-existing for *WShift*, suggesting that subjects with a heuristic reasoning are less committed to a specific course of action. We also notice that there are significantly more subjects who do not play *O1* but play *O2* (left column without bottom cell) than subjects who do not play *O2* but play *O1* (bottom row without leftmost cell). The results are reinforced when we condition on the choice in task 1 (right graph). With the already mentioned exception of subjects who play *WShift*, the best predictor for an individual’s strategy in task 2 is his strategy in task 1.

To further analyze individual choices across tasks, we perform a similar Probit regression as in Table 5, where the dependent variable is whether the individual follows strategy *O2* in task 2 (regressions (1), (2), (3) - *O2*) and whether he follows strategy *R* in task 2 (regressions (4), (5), (6) - *R*). For the regressors, we include dummies for all strategy choices in task 1 except *Unclassified* (*O1*, *WStay*, *WShift*, *R*), dummies for gender and *3same* in task 2, as well as dummies for the age-categories **C2** and above. The results are collected in Table 6.

	<i>O2</i> (in task 2)			<i>R</i> (in task 2)		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>O1</i> (in task 1)	0.999** (0.221)	0.973** (0.224)	0.510* (0.251)	-0.413 (0.268)	-0.611* (0.281)	-0.272 (0.313)
<i>WStay</i> (in task 1)	0.062 (0.296)	0.040 (0.295)	0.068 (0.320)	-0.616 (0.387)	-0.677 (0.400)	-0.687 (0.430)
<i>WShift</i> (in task 1)	0.062 (0.201)	0.060 (0.202)	-0.080 (0.216)	0.108 (0.206)	0.093 (0.212)	0.198 (0.220)
<i>R</i> (in task 1)	-0.099 (0.199)	-0.139 (0.204)	-0.232 (0.222)	0.629** (0.193)	0.521** (0.199)	0.479* (0.212)
<i>female</i>	—	0.139 (0.141)	-0.002 (0.152)	—	-0.403** (0.151)	-0.303 (0.159)
<i>3same</i> (in task 2)	—	-0.144 (0.169)	0.204 (0.184)	—	-0.693** (0.199)	-0.884** (0.207)
C2	—	—	-0.623* (0.294)	—	—	0.767** (0.246)
C3	—	—	0.380 (0.219)	—	—	0.322 (0.235)
C4	—	—	0.795** (0.248)	—	—	-0.096 (0.277)
C5	—	—	1.409** (0.280)	—	—	-1.177* (0.474)
<i>constant</i>	-0.736** (0.145)	-0.758** (0.174)	-1.025** (0.243)	-0.849** (0.150)	-0.441* (0.179)	-0.696** (0.245)
Observations	382	382	382	382	382	382
AIC	435.3	437.6	391.9	404.6	389.5	364.4

* $p < 0.05$; ** $p < 0.01$; (st. errors in parenthesis)

Table 6: Probit regressions of strategies in task 2 as a function of strategies in task 1.

Confirming the highlighted correlation of strategies across tasks, we notice that playing $O1$ is predictive of playing $O2$ whereas none of the other strategies is. Similarly, R in task 1 is also predictive of R in task 2. Subjects who play the optimal strategy in task 1 are also less likely to play the reverse strategy in task 2, but this effect disappears when we control for the age-category. Finally, the effects of age and of three identical draws emphasized in Table 5 persist after we control for the strategy in task 1.

4 Concluding remarks

In our experiment, children and teenagers learn the composition of balls in an urn through sampling with replacement. The design allows us to study learning strategies across ages. In particular, it makes it possible to assess the developmental trajectory of inference-based strategies (including bayesian) and mechanistic learning. Relevant comparisons across age groups are promoted by our unique population of subjects who all attend the same school, are subject to the same curriculum and are homogenous in terms of socioeconomic status. This strength comes at the cost of a smaller sample size per age group compared to standard experimental economics studies in adults. Also, notwithstanding that our adult control population is far from a perfect match for that population, it still provides a comparison group useful to establish baseline behavior in the exact same task.

The strong departures from optimal bayesian behavior observed in our experiment are consistent with the anomalies reported in adult learning. Wrong inferences and mechanistic rules are pervasive in our study across all ages. However, our study unveils an important developmental trajectory. Mechanistic rules are more frequent in young children and their usage decreases with age. Over time, they are replaced first by wrong inferences and then by optimal bayesian behavior. This means that some of the biases observed in young adults are not acquired through time to simplify problems. Instead, they are in place early in life and, if anything, tend to be replaced by sophisticated bayesian inferences. Naturally, the change is not unidirectional and, as reported in the psychology and behavioral economics literature, other biases develop over the years as a response to personal experience and environmental circumstances.

A possible explanation for the coexistence of sophisticated and mechanistic abilities is that, as for other decision problems, learning may be achieved via either goal-directed or habitual processes, which are self-selected depending on the learning context. Since goal-directed mechanisms crucially depend on cognitive functions, it is plausible that sophisticated bayesian learning is associated with cognitive maturity. As such, humans may rely on mechanistic rules easy to implement by default and switch to a cognitively intense bayesian method whenever appropriate. For young adults, whose cognitive abilities and

working memory are at the peak, the trade-off between bayesian inferences and heuristic rules is favorable to the former. By contrast, for children and older adults the trade-off is favorable to the latter. Such idea would be consistent with the developmental trajectory emphasized in the present study, as well as the results on heuristic decision making by older adults emphasized in the literature (Besedeš et al., 2012; Brocas et al., 2015).

References

- Jonathan Baron, Laura Granato, Mark Spranca, and Eva Teubal. Decision-making biases in children and early adolescents: Exploratory studies. *Merrill-Palmer Quarterly (1982-)*, pages 22–46, 1993.
- Jeffrey M Beck, Wei Ji Ma, Roozbeh Kiani, Tim Hanks, Anne K Churchland, Jamie Roitman, Michael N Shadlen, Peter E Latham, and Alexandre Pouget. Probabilistic population codes for bayesian decision making. *Neuron*, 60(6):1142–1152, 2008.
- Tibor Besedeš, Cary Deck, Sudipta Sarangi, and Mikhael Shor. Age Effects and Heuristics in Decision Making. *Review of Economics and Statistics*, 94(2):580–595, 2012.
- Elizabeth Bonawitz, Stephanie Denison, Alison Gopnik, and Thomas L Griffiths. Win-stay, lose-sample: A simple sequential algorithm for approximating bayesian inference. *Cognitive psychology*, 74:35–65, 2014.
- Fatemeh Borhani and Edward J Green. Identifying the occurrence or non occurrence of cognitive bias in situations resembling the monty hall problem. *arXiv preprint arXiv:1802.08935*, 2018.
- Isabelle Brocas, Juan Carrillo, T Dalton Combs, and Niree Kodaverdian. Consistency in simple vs. complex choices over the life cycle. Working Paper, 2015.
- Isabelle Brocas, Juan D Carrillo, and Niree Kodaverdian. Altruism and strategic giving in children and adolescents. Working Paper, 2017.
- Colin Camerer and Teck Hua Ho. Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874, 1999.
- Gary Charness, Uri Gneezy, and Brianna Halladay. Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization*, 131:141–150, 2016.
- Rachel Croson and James Sundali. The gambler’s fallacy and the hot hand: Empirical data from casinos. *Journal of risk and uncertainty*, 30(3):195–209, 2005.
- Denise Davidson. The representativeness heuristic and the conjunction fallacy effect in children’s decision making. *Merrill-Palmer Quarterly*, pages 328–346, 1995.
- Sophie Deneve. Bayesian spiking neurons 1: inference. *Neural computation*, 20(1):91–117, 2008.

- Stephanie Denison, Vashti Garcia, Kathleen Konopczynski, and Fei Xu. Probabilistic reasoning in preschoolers: random sampling and base rate. In *Proceedings of the Cognitive Science Society*, volume 28, 2006.
- Ido Erev and Alvin E Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, pages 848–881, 1998.
- Kurt W Fischer and Louise Silvern. Stages and individual differences in cognitive development. *Annual Review of Psychology*, 36(1):613–648, 1985.
- Thomas Gilovich, Robert Vallone, and Amos Tversky. The hot hand in basketball: On the misperception of random sequences. *Cognitive psychology*, 17(3):295–314, 1985.
- Jacqueline J Goodnow and Thomas F Pettigrew. Effect of prior patterns of experience upon strategies and learning sets. *Journal of Experimental Psychology*, 49(6):381, 1955.
- Alison Gopnik. Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337(6102):1623–1627, 2012.
- Alison Gopnik and Joshua B Tenenbaum. Bayesian networks, bayesian learning and cognitive development. *Developmental science*, 10(3):281–287, 2007.
- David M Grether. Testing bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior & Organization*, 17(1):31–57, 1992.
- Charles A Holt and Angela M Smith. An update on bayesian updating. *Journal of Economic Behavior & Organization*, 69(2):125–134, 2009.
- Janis E Jacobs and Maria Potenza. The use of judgement heuristics to make social and object decisions: A developmental perspective. *Child Development*, 62(1):166–178, 1991.
- Daniel Kahneman and Amos Tversky. Subjective probability: A judgment of representativeness. *Cognitive psychology*, 3(3):430–454, 1972.
- Shulamith Kreitler and Hans Kreitler. Development of probability thinking in children 5 to 12 years old. *Cognitive Development*, 1(4):365–390, 1986.
- Wei Ji Ma, Jeffrey M Beck, and Alexandre Pouget. Spiking networks for bayesian inference and choice. *Current opinion in neurobiology*, 18(2):217–222, 2008.
- Joshua Benjamin Miller and Adam Sanjurjo. Surprised by the gambler’s and hot hand fallacies? a truth in the law of small numbers. Working Paper, 2016.

- Zoi Nikiforidou and Jenny Pange. The notions of chance and probabilities in preschoolers. *Early childhood education journal*, 38(4):305–311, 2010.
- Jean Piaget. *The theory of stages in cognitive development*. McGraw-Hill, 1971.
- Jean Piaget and Bärbel Inhelder. *La genèse de l'idée de hasard chez l'enfant*. Presses universitaires de France, 1951.
- Chrysanthi Skoumpourdi, Sonia Kafoussi, and Konstantinos Tatsis. Designing probabilistic tasks for kindergartners. *Journal of early childhood research*, 7(2):153–172, 2009.
- Mark Steyvers, Michael D Lee, and Eric-Jan Wagenmakers. A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3):168–179, 2009.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press, 2011.
- Dek Terrell. A test of the gambler's fallacy: Evidence from pari-mutuel games. *Journal of risk and uncertainty*, 8(3):309–317, 1994.
- Edward Lee Thorndike. *Animal intelligence: Experimental studies*. Macmillan, 1911.
- Amos Tversky and Daniel Kahneman. Belief in the law of small numbers. *Psychological bulletin*, 76(2):105, 1971.
- Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- Fei Xu and Joshua B Tenenbaum. Word learning as bayesian inference. *Psychological review*, 114(2):245, 2007.

Appendix A. Instructions


A1. Instructions of Task 1.

Here's how the game will go. We have a box here and inside there are 12 balls, 6 green balls and 6 yellow balls. [Empty contents of box onto table so everyone sees]

HELPER1 will take one ball out without looking and he/she will not look at it either. Please pick a ball HELPER1 and hide it in your pocket. [HELPER1 removes one ball and hides it in their pocket].

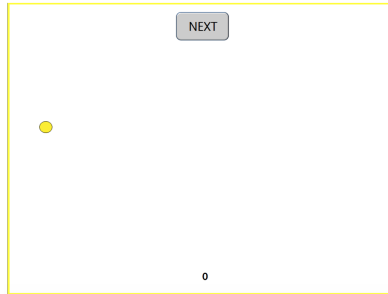
Now there are 11 balls in the box. Can we know the colors of the remaining 11 balls? Let's see. There were 12 balls, 6 green and 6 yellow and HELPER1 took one.

If HELPER1 removed a yellow ball, there are 6 green balls and 5 yellow balls in this box. But if HELPER1 removed a green ball, there are 5 green balls and 6 yellow balls in this box. Is everybody following? So, there are 5 green balls for sure and there are 5 yellow balls for sure, but the last ball could be either green or yellow.



Let's start with a pretend game. HELPER1 is going to pick a ball from the box and show everyone. When we play the game for real, one of you will pick the ball and will have to close the eyes to pick the ball. [HELPER1 chooses a yellow ball from the box and shows everyone].

Your screen would look like this:




Now HELPER1 puts it back in the box so there are still 11 balls in there. [HELPER1 returns ball to box].

HELPER1, go ahead and pick a second ball. [HELPER1 chooses a green ball from the box and shows everyone, then returns it to the box].

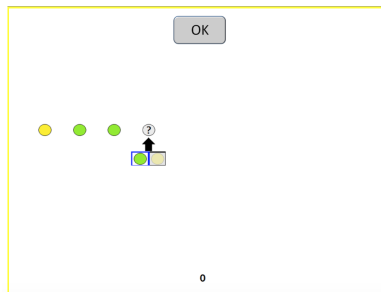
So now this is what you would see on your screen:



HELPER1 will pick another ball and show everyone then put it back in the box. [HELPER1 chooses a green ball from the box and shows everyone, then returns it to the box]. This is what you would see on your screen:

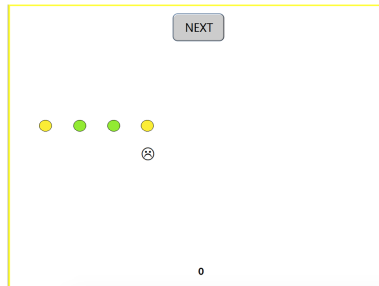


After HELPER1 has picked three balls, it's time to guess what he/she will draw next. Remember to guess quietly and not blurt out your guess! You will see a screen like this:

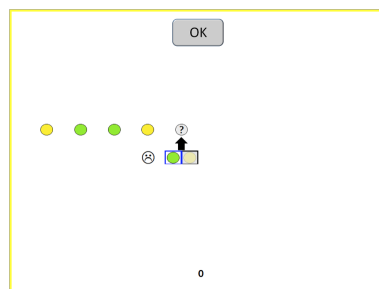


You will see a question mark and a yellow and green ball underneath that. You will tap on the color you think HELPER1 will draw next. I'm going to go ahead and tap on the green ball, then click the "OK" button to lock in my guess. Remember, in the actual game you will think about your guess and quietly submit it, rather than say it out loud.

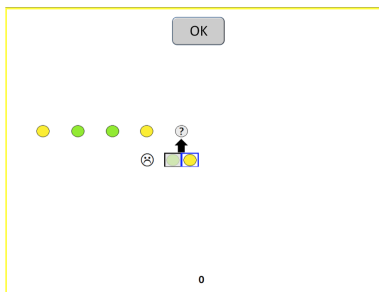
Go ahead and pick out another ball, HELPER1. [HELPER1 chooses a yellow ball].
Awww, too bad. OK, so this is what my screen would look like:



The fourth ball is yellow because HELPER1 picked a yellow ball. There is a sad face under it because my guess was wrong - I guessed that he/she would pick a green ball.
Let's do one more for pretend. I want to guess green. My screen looks like this now:



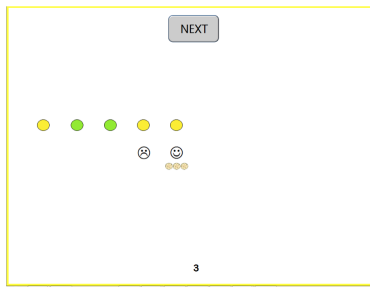
Actually, I want to guess yellow.



See, you can change your guess but once you press "OK" your guess is locked in.
HELPER1, please pick a ball. [HELPER1 draws a yellow ball].

It was yellow! My guess was right and here's what I would see on my screen:

The fifth ball is yellow because HELPER1 picked a yellow ball. There is a happy face



under it because I guessed right. Every time you guess correctly, you will win 3 tokens. On the very bottom, you'll see how many tokens you earned so far. Remember, more tokens means more money on your amazon gift card (or toys, in the case of younger cohorts) so think carefully.


Are there any questions so far?

Let's play for real now. HELPER1, please return the hidden ball to the box. Now we have again 6 yellow and 6 green balls.

Remember everyone, do not say what you think out loud and do not look at anyone's screen but your own. There is no talking aloud.

OK HELPER1, please pick a ball without looking and hide it in your pocket. [HELPER1 removes one ball and hides it in another box or bag].

Again, we have either 6 green balls and 5 yellow balls or 5 green balls and 6 yellow balls.



Now, I will ask one of you to pick a ball without looking and show everyone, then put it back in the box. You can all start thinking about what will be picked next but please remain silent. [Repeat 2 times]

Time to guess. Please quietly tap the color you think will be picked next on your tablet. You can change your mind but when you press "OK" your guess will be locked-in. Everyone made a guess? [Ask a child to pick a ball, show everyone, and return it to the box].

[Display draw on tablets, provide feedback, advance to next guess] [Repeat 6 times]

Let's see the hidden ball now, HELPER1! And it was (green/yellow)!


A2. Instructions of Task 2.

We're going to play this game one more time but there will be one difference. Right now, there are 12 balls, 6 green and 6 yellow in the box. I need you to pick one ball HELPER1 and show everyone. [HELPER1 picks a yellow or green ball and shows everyone]

If a yellow ball was removed. So now there are 11 balls left in the box, 6 green and 5 yellow. We'll do the same thing as before from here on. HELPER1 will take one ball out without looking and he/she will not look at it either.

Please pick a ball HELPER1 and hide it in this box/bag. [HELPER1 removes one ball and hides it in another box or bag].


Now there are 10 balls in the box. Can we know the colors of the remaining 10 balls? Let's see. There were 11 balls, 6 green and 5 yellow and HELPER1 took one. If HELPER1 removed a yellow ball, there are 6 green balls and 4 yellow balls in this box. But if HELPER1 removed a green ball, there are 5 green balls and 5 yellow balls in this box. Is everybody following? So, there are 4 green balls for sure and there are 4 yellow balls for sure, but the last 2 balls could either both be green or one could be green and one could be yellow. The box looks like one of these two:



If a green ball was removed. So now there are 11 balls left in the box, 5 green and 6 yellow. We'll do the same thing as before from here on. HELPER1 will take one ball out without looking and he/she will not look at it either.

Please pick a ball HELPER1 and hide it in this box/bag. [HELPER1 removes one ball and hides it in another box or bag].

Now there are 10 balls in the box. Can we know the colors of the remaining 10 balls? Let's see. There were 11 balls, 5 green and 6 yellow and HELPER1 took one. If HELPER1 removed a yellow ball, there are 5 green balls and 5 yellow balls in this box. But if HELPER1 removed a green ball, there are 4 green balls and 6 yellow balls in this box. Is everybody following? So, there are 4 green balls for sure and there are 4 yellow balls for sure, but the last 2 balls could either both be yellow or one could be green and one could be yellow. The box looks like one of these two:



Now, one of you will pick the first ball without looking, show everyone and put it back in the box. You can start thinking about what will be picked next but please remain silent. [Repeat 2 times]

Time to guess. Please quietly tap the color you think will be picked next on your tablet. You can change your mind but when you press “OK” your guess will be locked-in. Everyone made a guess? [Ask a child to pick a ball, show everyone, then return it to the box].

[Display draw on tablets, provide feedback, advance to next guess] [Repeat 6 times]
 Let’s see the hidden ball now, HELPER1! And it was (green/yellow)!

Appendix B. Equilibrium choice by grade

We present the same information as in Fig.2 (left), that is, the average proportion of bayesian choices in tasks 1 and 2. However, we separate children by grades (kindergarten to 11) instead of grouping them in age-categories with three grades each. The results are compiled in Fig.9.

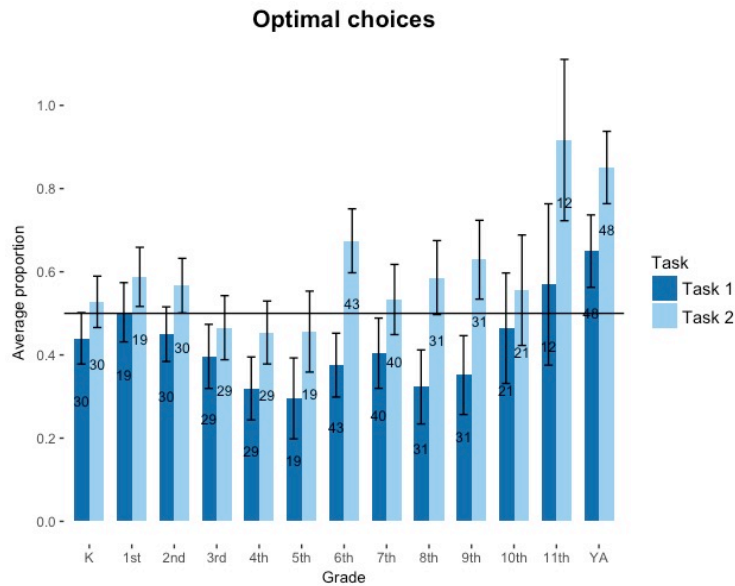


Figure 9: Proportion of bayesian choices in task 1 and task 2 by grade.