Novelty and uncertainty differentially drive exploration across development

Kate Nussenbaum*[1], Rebecca E. Martin*[1,2], Sean Maulhardt[1,3], Yi (Jen) Yang[1,4], Greer Bizzell-Hatcher[1], Naiti S. Bhatt[1], Maximilian Scheuplein[1,5], Gail M. Rosenbaum[1,6], John P. O'Doherty[7], Jeffrey Cockburn[7], Catherine A. Hartley[1]

** Shared first authorship*

1. New York University
2. University of Pennsylvania
3. University of Maryland
4. Temple University
5. Leiden University
6. Geisinger Health System
7. Caltech

Abstract

Across the lifespan, individuals frequently choose between exploiting known rewarding options or exploring unknown alternatives. A large body of work has suggested that children may explore more than adults. However, because novelty and reward uncertainty are often correlated, it is unclear how they differentially influence decision making across development. Here, children, adolescents, and adults (ages 8 - 27 years, N = 122) completed an adapted version of a recently developed value-guided decision-making task (Cockburn et al., 2021) that decouples novelty and uncertainty. In line with prior studies, we found that exploration decreased with increasing age. Critically, participants of all ages demonstrated a similar bias to select choice options with greater novelty, whereas aversion to reward uncertainty increased into adulthood. Computational modeling of participant choices revealed that whereas adolescents and adults demonstrated attenuated uncertainty aversion for more novel choice options, children did not factor reward uncertainty into their decisions at all.

Across the lifespan, exploration increases individuals' knowledge of the world and promotes the discovery of rewarding actions. In some circumstances, exploring new options may yield greater benefits than sticking to known alternatives, whereas in others, "exploiting" known options may bring about greater rewards. This trade-off is known as the 'explore-exploit' dilemma (Cohen et al., 2007; Sutton et al., 1998), reflecting the challenge inherent to resolving this tension. In general, the optimal balance between exploration and exploitation may shift across the lifespan. Relative to adults, children tend to know less about the world and have longer temporal horizons over which to exploit newly discovered information (Gopnik, 2020; Gopnik et al., 2017). Thus, it may be advantageous to explore to a greater extent earlier in life, and gradually shift to a more exploitative decision strategy as experience yields knowledge. Empirical data suggest that individuals at varied developmental stages do indeed tackle explore-exploit problems differently. Children and adolescents tend to explore more than adults (Christakou et al., 2013; Giron et al., 2022; Jepma et al., 2020; Lloyd et al., 2020; Nussenbaum & Hartley, 2019; E. Schulz et al., 2019), and this increased exploration promotes enhanced learning about the structure of the environment (Blanco & Sloutsky, 2020; Liquin & Gopnik, 2022; Sumner et al., 2019). Despite compelling arguments for why an early bias toward exploration may be advantageous and growing evidence that children *are* in fact more exploratory than adults, the cause of the developmental shift toward exploitation remains unclear.

Prior work has revealed that across development, two features of choice options influence exploration: stimulus *novelty* (Daffner et al., 1998; Gottlieb et al., 2013; Henderson & Moore, 1980; Jaegle et al., 2019; Kakade & Dayan, 2002; Wittmann et al., 2008) and reward *uncertainty* (Badre et al., 2012; Blanco & Sloutsky, 2020; Gershman, 2018; Somerville et al., 2017; Trudel et al., 2021; Wang et al., 2021; Wilson et al., 2014). Here, we use *novelty* to refer to the extent to which choice options have been previously encountered and *uncertainty* to refer to the variance in the distributions of rewards that they yield. Disentangling the role of novelty and uncertainty in driving exploratory decision making is challenging because they are often correlated. For example, a new toy has high novelty because it has never been encountered and high reward uncertainty because its entertainment value is unknown. Still, while novel stimuli almost always have high reward uncertainty, in many cases, *familiar* options do as well — when buying a familiar toy as a gift for *someone else*, one may have little knowledge of how much they will like it.

A recent study in adults took advantage of these types of choices, and, by harnessing familiar options with unknown reward probabilities, decoupled the influence of novelty and uncertainty on exploratory decision-making in adults (Cockburn et al., 2021). Adults were novelty-seeking, preferentially selecting choice options that they had encountered infrequently in the past versus those that were more familiar. However, adults were also uncertainty averse, such that they tended to avoid options with high reward uncertainty. This tension between avoiding uncertain options while pursuing novel alternatives, which are themselves inherently uncertain, suggested *interactive* effects of choice features. Computational modeling further revealed that stimulus novelty diminished the influence of uncertainty on exploratory choice. Thus, these findings suggest that value-guided decision making in adults — and specifically, the balance between exploration and exploitation — may be governed by complex interactions among different features of choice options. To date, however, the influences of novelty and uncertainty have not been disentangled in children and adolescents.

Changes in the influence of these choice features may shift the explore-exploit balance across development. A stronger appetitive influence of stimulus novelty may drive greater exploration earlier in life. Reduced uncertainty aversion, or perhaps even an early preference to explore more uncertain options may similarly promote heightened exploratory behavior. Novelty and uncertainty may also exert unique, interactive effects for younger individuals. Though prior studies have found effects of novelty (Henderson & Moore, 1980; Mendel, 1965) and reward uncertainty (Blanco & Sloutsky, 2020; Meder et al., 2021; E. Schulz et al., 2019) on exploration and choice in early childhood, no prior studies have charted how their influence changes from childhood to early adulthood, leaving open the question of *why* children tend to explore more than adults. Further, in most developmental studies, novelty and uncertainty are confounded, making it difficult to tease apart their separate, motivational effects.

Here, using an adapted version of the task introduced in Cockburn et al. (2021) with a large age-continuous developmental sample, we asked how the influence of novelty and uncertainty on exploratory choice changes from middle childhood to early adulthood.We hypothesized that the developmental shift from more exploratory to more exploitative behavior would be driven by changes in how both novelty and reward uncertainty affect the evaluation of choice options from childhood to adulthood.

Methods

Participants

One hundred and twenty two participants between the ages of 8 and 27 years old (Mean age = 17.9 years, SD age = 5.6 years, 62 female, 59 male, 1 non-binary) completed the experiment. Based on prior, similar studies of value-guided decision-making from childhood to adulthood (Habicht et al., 2021; Somerville et al., 2017), we determined a target sample size of *N*

= 120, evenly distributed across our age range, prior to data collection. The final analyzed sample of 122 participants comprised *n* = 30 children (mean age = 10.5 years; range = 8.1 - 12.7 years; 14 females), *n* = 30 adolescents (mean age = 15.5 years; range = 13.6 - 17.8 years; 16 females), and *n* = 62 adults (mean age = 22.6 years; range = 18.1 - 27.8 years; 32 females). Data from one additional participant was not analyzed because the participant chose to stop the experiment prior to completing the entire exploration task. Two participants included in the final analyzed sample were excluded from memory test analyses due to technical errors during data acquisition.

Participants were recruited from the local New York City community. Participants reported normal or corrected-to-normal vision and no history of diagnosed psychiatric or learning disorders. Based on self- or parental report, 33.6% of participants were Asian, 33.6% were White, 18.0% were two or more races, 13.1% were Black, and 1.6% were Pacific Islander / Native Hawaiian. In addition, 15.6% of participants were Hispanic.

Task

Exploration task. Participants completed a child-friendly decision-making task adapted from one used in a prior adult study (Cockburn et al., 2021). The child-friendly version of the task was framed within an "Enchanted Kingdom" narrative and included fewer stimuli and trials per block than the version used in prior work (Cockburn et al., 2021) with adults. Within this narrative framework, participants were tasked with finding gold coins to raise money to build a bridge to unite the two sides of the kingdom. Various creatures had hidden the gold coins in different territories around the kingdom. On every trial, participants had to choose between two hiding spots to search for a coin.

The task was divided into 10 blocks of 15 trials. Each block took place within a different territory in which coins were hidden by a distinct creature. Each creature hid their coins among three possible locations. Half of the blocks were 'easy' — in easy blocks, the creature's favorite hiding spot held a coin on 80% of trials, their second favorite held a coin on 50% of trials, and their least favorite held a coin on 20% of trials. The other half of blocks were 'hard' — in hard blocks, the creature's favorite hiding spot held a coin on 70% of trials, their second favorite held a coin on 50% of trials, and their least favorite held a coin on 30% of trials. Participants were not explicitly informed of these probabilities, and had to learn, through trial and error, where each creature was most likely to hide a coin.

On every trial, participants viewed two hiding spots and had 4 seconds to select one in which to search for a coin by pressing one of two keys on a standard keyboard (Figure 1). After a brief delay in which the option they selected was outlined (500ms), participants saw the outcome of their choice — either a coin or an X indicating that they had not found a coin (1.5 seconds). Throughout each block, the background of the screen indicated the territory and a picture in the lower left corner indicated the creature that had hidden coins there.

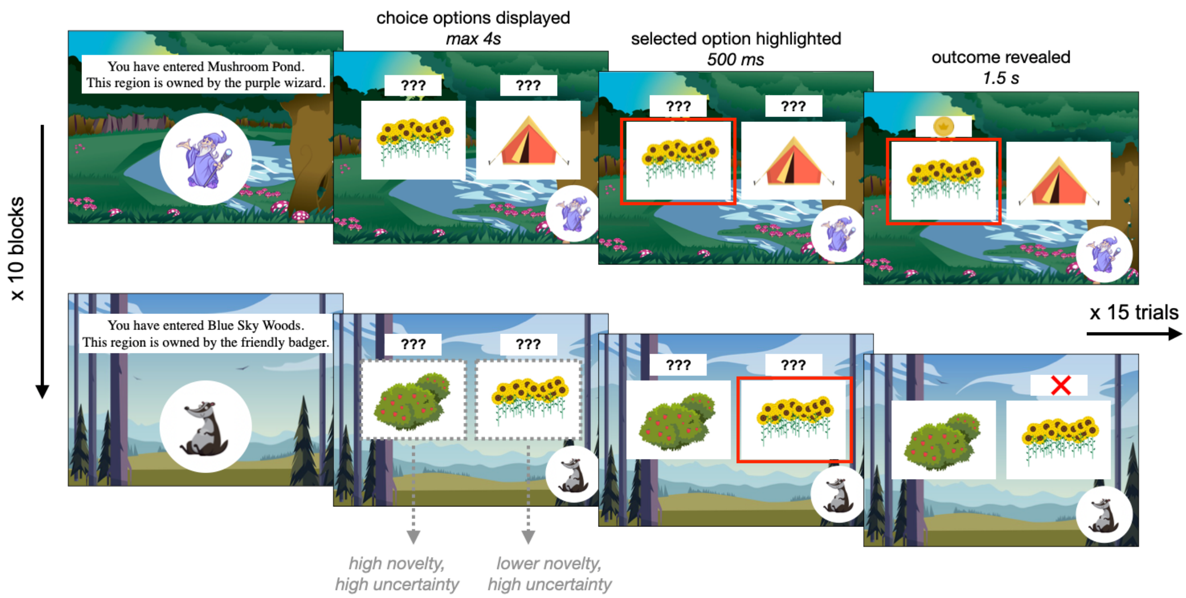

Figure 1: Exploration task. Participants completed 10 blocks of 15 choice trials in which they selected between two of three "hiding spots" to find gold coins. Within each block, two hiding spots had been previously encountered and one was completely novel. Each block took place within a different 'territory' in which a new creature hid coins. Each creature had different preferred hiding spots, such that the reward probabilities associated with each option were reset at the beginning of each block.

Importantly, after the first block, each subsequent block contained two hiding spots that participants had encountered in previous blocks and one novel hiding spot they had not seen before. Though participants had already encountered two of the hiding spots within each block, the reward probabilities were re-randomized for every creature. In this way, the task dissociated sensory *novelty* and reward *uncertainty*. At the beginning of every block, the novelty of each hiding spot varied — at least one hiding spot was completely novel, whereas from the second block on, the other two had been encountered anywhere from 4 to 82 times (mean = 22.4 encounters) — but *all* hiding spots had high reward uncertainty. Participants were explicitly told that the reward probabilities were reset in every block; within the task narrative, this was framed as each creature having different favorite hiding spots in their respective territory (see supplement for analyses demonstrating that participants of all ages indeed comprehended these instructions and 'reset' the reward probabilities at the beginning of each block).

The order of the creatures, the hiding spots assigned to each creature, and the order of easy and hard blocks were randomized for each participant. Within each block, the reward probabilities assigned to the two 'old' hiding spots and the novel hiding spot were randomized. On each trial, the two hiding spots that appeared as choice options and their positions on the screen (left or right) were also randomized

Memory test. Immediately after the exploration task, participants completed a surprise memory test. They were shown each of the ten creatures, one at a time, and asked to select its favorite hiding spot from an array of five options, using numbers on the keyboard. The array of five options always included the hiding spot in which the creature was most likely to hide the coin (the correct answer), the two other hiding spots where that creature hid coins, a previously encountered hiding spot from a different block of the task, and a new hiding spot that was not presented in the exploration task.

Instructions and practice. Prior to completing both the exploration and the memory task, participants went through extensive, interactive instructions with an experimenter. The instructions were written and illustrated on the computer screen, and an experimenter read them aloud. During the instructions, participants were informed that a.) hiding spots may repeat throughout the task but each creature had different favorite hiding spots, b.) a creature would not *always* hide its coins in its favorite spot, and c.) each creature's hiding spot preferences remained stable throughout the entire block. They also went through three comprehension questions with an experimenter and completed two full practice blocks with stimuli that were not used in the main task.

WASI

After the exploration task and memory test, participants were administered the vocabulary and matrix reasoning subtests of the Wechsler Abbreviated Scale of Intelligence (Wechsler, 2011). Because our primary aim was to address the relation between age and exploratory behavior, we report results from the models *without* WASI scores in the main text of the manuscript and include results from the models with WASI scores in the supplement.

Analysis approach

Treatment of age. We treated age as a continuous variable in all regression analyses. We binned participants into three age groups (children aged 8 - 12 years, adolescents aged 13 - 17 years, and adults aged 18 - 27 years) for data visualization and model comparison purposes. Our adult age group spanned a greater age range and included double the number of participants as our child and adolescent age groups; as such, we include additional visualizations and model comparison results that subdivide this group into college-aged adults (18 - 22 years) and post-college-aged adults (23 - 27 years) in the supplement.

Because we originally hypothesized that the influence of expected value, uncertainty, and novelty would change monotonically with age, we included linear age in all regression analyses. To account for potential non-linearities in our age effects, we report results from analyses including quadratic age in the supplement, though we note that the quadratic age effects we observed do not hold when we control for potential cohort-level IQ differences across our sample (see supplement).

Defining expected value, uncertainty, and novelty. To examine the influence of expected value, uncertainty, and stimulus novelty on choice behavior, we defined and computed these three feature values for each choice option on every trial (Cockburn et al., 2021). Expected value was defined as the mean of the beta distribution specified according to the win and loss history of each choice option (hiding spot) within the block: $\frac{\alpha}{\alpha + \beta}$ where α = number of wins + 1 and β = number of losses + 1. Uncertainty was defined as the variance of this beta distribution: $\frac{(\alpha * \beta)^2}{\alpha + \beta + 1}$. Stimulus novelty was determined by taking the variance of a different beta distribution, where α = the number of times participants had seen the choice option before throughout the entire task + 1 and β = 1.

Computational modeling. To more precisely quantify how learned value, uncertainty, and novelty influenced choice across age, we fit participant choice data with six different reinforcement learning models (see supplement). Across models, we conceptualized the learning process as that of a 'forgetful Bayesian learner,' such that the expected value of each choice option is computed as the mean of a beta distribution with hyperparameters that reflect recency-weighted win and loss outcomes (Cockburn et al., 2021). We then modified this baseline model by adding either fully separable or interacting uncertainty and novelty biases. Specifically, beyond the baseline model, we fit three additional models in which uncertainty and novelty exerted separable influences on choice behavior: a model augmented with a *novelty bias* that adjusted the initial hyperparameters of each option's beta distribution, a model augmented with an *uncertainty bias* that added or subtracted each option's scaled uncertainty to its expected utility, and a model augmented with both biases.

We additionally fit two models that account for interactions between novelty and uncertainty. Given the findings of Cockburn et al. (2021), we hypothesized that novelty may buffer the aversive influence of reward uncertainty. In other words, we expected that the extent to which the uncertainty of a given choice option would influence its utility would increase in relation to its familiarity. Thus, we fit two additional models (with and without a separate novelty bias) in which the uncertainty bias was ‘gated’ by stimulus familiarity (though the model with both a novelty bias and familiarity gate was not recoverable; see ‘methods’).

To test for age-related change in the way that novelty and uncertainty influence exploratory choice, we compared model fits for these six models within each age group using a random-effects Bayesian model selection procedure with simultaneous hierarchical parameter estimation (Piray et al., 2019). This model selection procedure estimates a population-level distribution of models and treats individual participants as random draws from this distribution, allowing different individuals within each age group to be best fit by different models. To determine the best-fitting model within each age group, we examined protected exceedance probabilities (PXPs), which reflect the probability that a given model in a comparison set is the most frequent, best-fitting model across participants, while controlling for differences in model frequencies that may arise due to chance.

Results

Exploration task performance

First, we examined whether participants learned to select the better options within each block. On each trial, the optimal choice was defined as the option with the higher reward probability. A mixed-effects logistic regression examining the effects of within-block trial number, age, block difficulty, and their interactions, with random participant intercepts and slopes across trial number and block difficulty, revealed that participants learned to make more optimal choices over the course of each block, $\chi^2(1)$ = 17.43, *p* < .001. In addition, participants made more optimal choices in easy relative to hard blocks, $\chi^2(1)$ = 24.9, *p* < .001 (Figure 2A). Finally, performance improved with increasing age, $\chi^2(1)$ = 5.44, *p* = .02 (Figure 2A). No interactions among trial number, age, and block difficulty reached significance (*p*s > .10).

We also examined how much reward participants earned throughout the task as a function of both age and block difficulty. If participants selected choice options at random, we would expect them to earn reward on 50% of the 75 “easy” trials and 50% of the 75 “hard” trials, for an average of 37.5 coins per block type. The majority of participants across our age range performed above chance (Figure 2B). Performance improved across age, with older participants earning significantly more reward than younger participants, β = .13, SE = .06, *t* = 2.11, *p* = .036. There was not a significant effect of block difficulty on reward earned, nor was there an age x block difficulty interaction effect (*p*s > .34).

Taken together, these findings indicate that participants across our age range learned to select rewarding choice options throughout each block, though the extent to which participants learned to 'exploit' the most rewarding choice options increased across age. Participants also demonstrated above-chance (defined as .2) memory for the most rewarding choice within each block (Mean = .25; SE = .01; $t(119) = 4.02$, $p < .001$). Memory accuracy did not vary across age, $\beta = -.046$, SE = .07, $z = -.69$, $p = .49$.

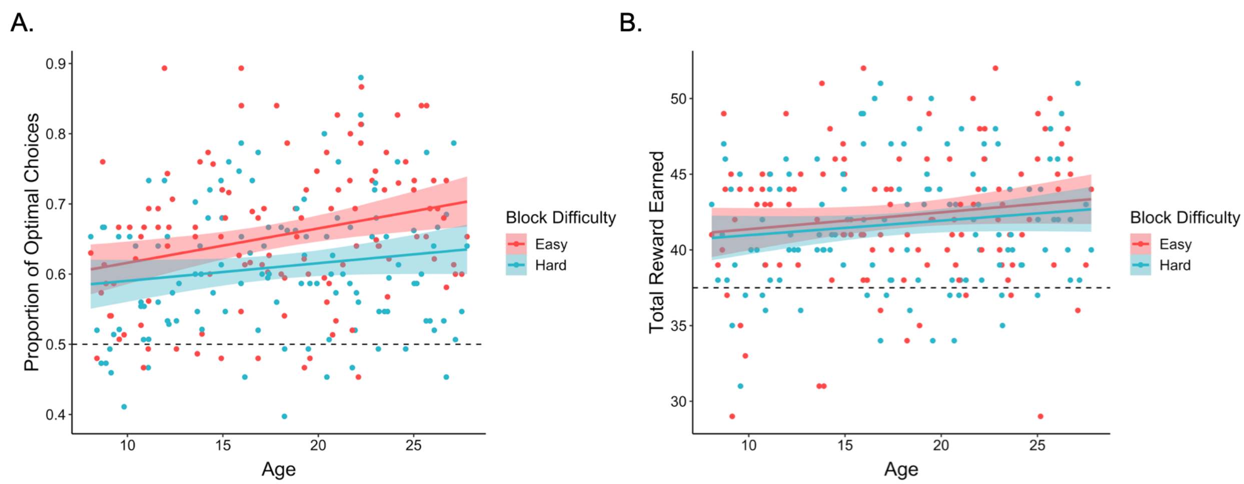

Figure 2. Exploration task performance. A) Participants made more optimal choices in the easy versus hard block of the task ($p < .001$). Choice performance also improved with increasing age ($p = .02$). B) Across blocks, older participants earned more reward than younger participants ($p = .036$). The lines show the best-fitting linear regression lines and the shaded regions around them represent 95% confidence intervals. The dotted lines indicate chance-level performance.

Age-related change in exploration

Next, we turned to our main questions of interest: whether and how novelty and uncertainty influenced exploration across age. To address this question, we computed the differences in expected value, uncertainty, and novelty between the left and right choice option on every trial. We then ran a mixed-effects logistic regression examining how these differences — as well as their interactions with age — related to the probability that the participant chose the left option on every trial. Participants were more likely to select the options that they had learned were more valuable, $\chi^2(1) = 155.5$, $p < .001$. However, expected value was not the only driver of choice behavior. Participants also demonstrated a bias toward selecting more novel stimuli, $\chi^2(1) = 105.72$, $p < .001$, and a bias *away* from choosing options with greater uncertainty, $\chi^2(1) = 14.1$, $p < .001$. On trials in which the two choice options had similar expected values (< .1 difference), participants selected the more novel option on 58.7% (SE = 0.8%) of trials and the more uncertain option on only 46.3% (SE = 0.8%) of trials.

The influence of expected value, novelty, and uncertainty on choice behavior each followed distinct developmental trajectories. Younger participants' choices were less value-driven relative to those of older participants, as reflected in a significant age x expected value interaction, $\chi^2(1)$ = 9.8, p = .002 (Figure 3A). Importantly, however, age-related increases in these 'exploitative' choices were *not* driven by age-related differences in novelty-seeking; there was not a significant interaction between age and novelty, $\chi^2(1)$ = 1.1, p = .301. In contrast to the relative stability of this novelty preference across age, we observed a significant age x uncertainty interaction effect, $\chi^2(1)$ = 12.2, p < .001, indicating greater uncertainty aversion in older participants (Figure 3A). All findings held when we included within-block trial number as an interacting fixed effect in the model (see supplement).

Corroborating these findings, on trials in which the two choice options had nearly identical expected values (< .1 difference), children, adolescents, and adults, on average, selected the more novel option on 59%, 59.2%, and 58.3% of trials, respectively (Figure 3B). However, whereas adults tended to avoid the more uncertain option, selecting it on only 42.1% of equal-expected-value trials, adolescents and children selected the more uncertain option on 48.9% and 52.5% of these trials respectively (Figure 3B). Thus, taken together, these results suggest that age-related decreases in exploratory choices were driven by an increase in aversion to reward uncertainty with increasing age.

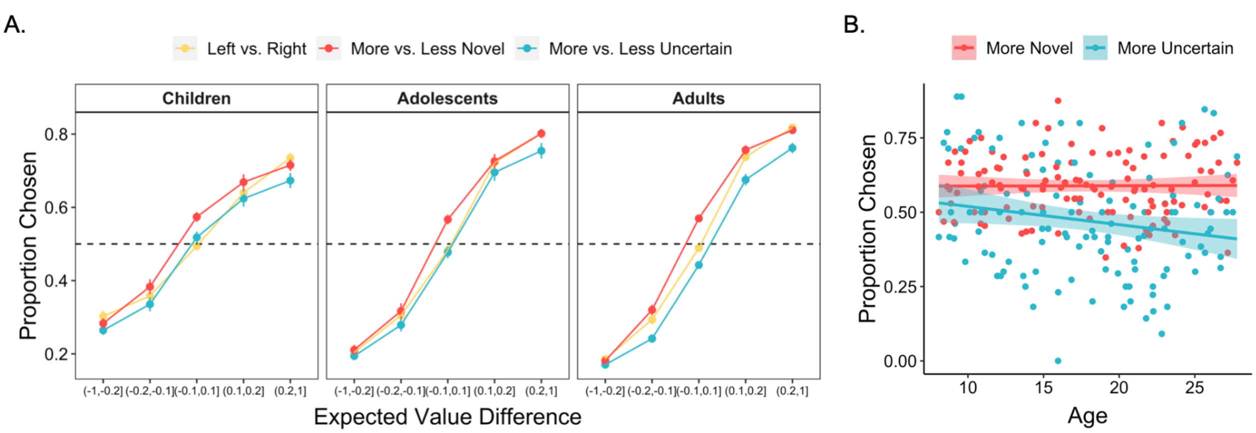

Figure 3. Influence of expected value, uncertainty, and novelty on choice behavior across age. A) The proportion of all trials in which the participants chose the left, more novel, and more uncertain choice option as a function of the expected value difference between the two options. Participants were more likely to choose options with greater expected value, higher novelty, and lower uncertainty (ps < .001). The influence of novelty did not vary across age, whereas uncertainty was more aversive in older participants (p < .001). B) The proportion of similar-expected-value trials (difference between the two options < .1) in which participants chose the more novel and more uncertain option, plotted as a function of continuous age. The lines show the best-fitting linear regression lines and the shaded regions around them represent 95% confidence intervals.

Computational characterization of choice

As in Cockburn et al. (2021), we observed opposing effects of novelty and uncertainty on choice — though participants sought out novel options, they shied away from those with greater uncertainty. At first glance, these results are somewhat puzzling because novel options are *inherently* uncertain. Reinforcement learning models that formalize different algorithms for how the expected utilities of the choice options are computed across trials can provide greater insight into how novelty and uncertainty may *interact* to influence exploratory choice behavior.

In line with findings from Cockburn et al. (2021), we found that adult choices were best characterized by a model in which choice utilities took into account *interactions* between novelty and uncertainty. Specifically, adult choices were best captured by the familiarity-gated uncertainty model (PXP Familiarity Gate = 1), in which uncertainty aversion was greater for more familiar options. Despite showing weaker aversion to uncertainty relative to adults, adolescents were also best fit by this model (PXP Familiarity Gate = 1). Children's choices, however, were best captured by a model with a simple novelty bias (PXP Novelty Bias = .62; PXP Familiarity Gate = .38).

Parameter estimates from the winning models reflected participants' bias toward novel stimuli and away from those with high reward uncertainty. Children's average "novelty bias" (from the group-level novelty bias model fits) was 1.49, indicating that they optimistically initiated the value of novel options. A one-sample t-test, implemented via the cbm model-fitting package (Piray et al., 2019), revealed that children's novelty bias parameter estimates were significantly different from 0, $t(16.3) = 5.96$, $p < .001$. The average value of the "uncertainty bias" (from the group-level familiarity-gated uncertainty model fits) was -.15 for both adolescents and adults. Uncertainty bias parameter estimates were significantly different from 0 in both age groups (Adolescents: $t(26.5) = -7.8$, $p < .001$; Adults: $t(11.3) = -4.16$, $p = .001$).

Model simulations revealed that the winning models well-captured qualitative features of behavioral choice data for each age group. For each model, we generated 50 simulated data sets using each of the 122 participants' trial sequence and parameter estimates (for a total of 6,100 simulated participants per model). Data from these simulations demonstrated that the familiarity-gated uncertainty model generated the most strongly diverging effects of novelty and uncertainty on choice, in line with the adult and adolescent data (Figure 4). The simpler novelty bias model instantiated a bias toward *both* novel and uncertain choices. Thus, these modeling results suggest that whereas adults and adolescents were more strongly deterred by the uncertainty of familiar options versus novel ones, children employed a simpler learning algorithm in which they optimistically initialized the value of novel choice options.

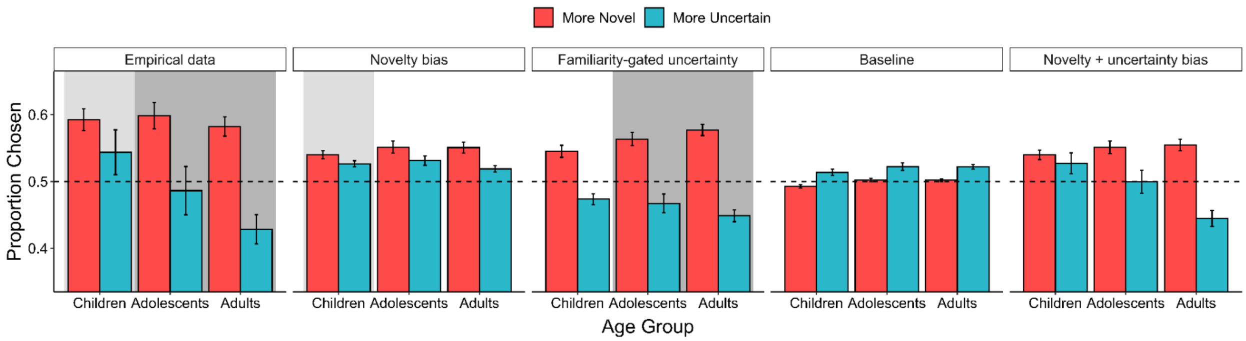

Figure 4. Model simulations. The average proportion of similar-expected-value trials (with expected value magnitude differences < .1) in which both real and simulated participants chose the more novel and more uncertain option. The shaded regions show the empirical data and best-fitting model for each age group. Error bars represented the standard error across participant means.

Discussion

In this study, we investigated how novelty and uncertainty influence exploration across development. Though new choice options tend to have *both* high novelty and high reward uncertainty, we found that the influence of these features on decision making follow distinct developmental trajectories. While participants across age demonstrated a similar bias toward selecting more novel choice options, only older participants showed aversion to those with greater uncertainty. These findings suggest that children's bias toward exploration over exploitation may arise from early indifference to reward uncertainty rather than heightened sensitivity to novelty.

Prior studies have found that novelty may be intrinsically rewarding (Wittmann et al., 2008), motivating individuals to approach, learn about, and remember new stimuli they encounter (Houillon et al., 2013; Krebs et al., 2009). Children (Henderson & Moore, 1980; Mendel, 1965; Valenti, 1985), adolescents (Spear, 2000), and adults (Cockburn et al., 2021; Daffner et al., 1998) all demonstrate novelty-seeking behavior. However, though many studies have shown novelty preferences at different developmental stages, little work has compared novelty preferences across age. Research in rodents has suggested that adolescents may demonstrate heightened sensitivity to novelty (Philpot & Wecker, 2008; Spear, 2000; Stansfield & Kirstein, 2006), but to the best of our knowledge, there is no human evidence for an adolescent peak in novelty preferences. Indeed, our findings suggest that an age-invariant drive to engage with novel stimuli promotes exploratory choice across development. Developmental differences in novelty-processing cannot fully account for developmental differences in exploration.

Whereas novelty-seeking did not exhibit age-related change, uncertainty aversion increased from childhood to early adulthood, potentially reflecting developmental improvement in the strategic modulation of information-seeking. Prior studies of decision making have found that individuals across age demonstrate uncertainty aversion in some environments (Camerer & Weber, 1992; Payzan-LeNestour et al., 2013; Rosenbaum & Hartley, 2019) and uncertainty-seeking in others (Blanchard & Gershman, 2018; Giron et al., 2022; E. Schulz et al., 2019). These seemingly discrepant patterns of behavior may be explained by differences in the utility of resolving uncertainty across contexts. In environments where learned information can be exploited to improve subsequent choices, resolving uncertainty has high utility (Rich & Gureckis, 2018; Wilson et al., 2014), whereas in choice contexts with short temporal horizons, there is little opportunity to use learned reward information to improve future decisions (Camerer & Weber, 1992; Levy et al., 2010). In our task, individuals had a relatively short horizon over which to exploit reward probabilities that themselves required multiple trials to learn — children's reduced uncertainty aversion may have emerged from insensitivity to the limited utility of gaining additional information about the most uncertain choice options (Somerville et al., 2017).

Another possibility is that the computational demands of tracking uncertainty may have precluded younger participants from using it to guide their decisions. Though some studies have found evidence for uncertainty-seeking in young children (Sumner et al., 2019), this behavior may arise incidentally from increased decision noise (Nussenbaum & Hartley, 2019) or from heuristic strategies like choice alternation. When children *do* purposefully direct their exploration to parts of the environment that are likely to yield the most information, uncertainty is often easy to discern. For example, in one study (Blanco & Sloutsky, 2020), children tended to select choice options with hidden reward amounts over those with visible reward amounts. In other studies with spatially correlated rewards, children could use the layout of revealed outcomes to direct their sampling toward unexplored regions (Giron et al., 2022; Meder et al., 2021; E. Schulz et al., 2019). In studies of causal learning and exploratory play, young children often use experiences of surprise or the coherence of their own beliefs about the world to direct their exploration toward uncertain parts of the environment (Bonawitz et al., 2012; L. E. Schulz & Bonawitz, 2007; Wang et al., 2021). These perceptual and metacognitive signals of uncertainty may be easier to evaluate than reward uncertainty in our task, which depended on distributions of binary outcomes. The developmental trajectory of uncertainty-seeking or uncertainty-averse behavior may be modulated by differences in the computational complexity of uncertainty estimation across environments.

Importantly, the use of stimulus *novelty* to guide decision-making likely relied on cognitive mechanisms distinct from those involved in the use of uncertainty. Whereas estimating stimulus novelty can be accomplished through recognition memory mechanisms, estimating reward uncertainty requires forming, retrieving, and in our task, overwriting, associations between choice options and the outcomes they yield. Binding information in memory and remembering items within specific contexts follow more protracted developmental trajectories than simply recognizing items as old or new (Lee et al., 2016; Shing et al., 2010, 2008). Given the high correlation between novelty and uncertainty in most environments, the use of novelty alone may promote exploration and the discovery of new information without requiring the complex computations and associative learning mechanisms needed to track uncertainty — mechanisms that may be particularly costly to engage early in life.

Interestingly, however, our computational modeling findings largely replicated previous work suggesting that by adulthood, novelty and uncertainty interact competitively (Cockburn et al., 2021), exerting opposing motivational influences on decision making. Using fMRI, this prior adult study (Cockburn et al., 2021) revealed that activation in the ventral striatum reflects a biased reward prediction error consistent with optimistic value initialization for novel stimuli, whereas activation in the medial prefrontal cortex (mPFC) reflects the subjective utility of uncertainty reduction. In line with choices being best characterized by a model in which the aversive influence of uncertainty was dampened by novelty, the integration of uncertainty into these mPFC value representations was similarly diminished for novel stimuli (Cockburn et al., 2021). Cockburn et al. (2021) posited that attenuated uncertainty aversion for novel stimuli may promote exploration even in environments where deriving the prospective utility of uncertainty reduction is difficult. Our findings further suggest that these competitive interactions between uncertainty and novelty may emerge in adolescence, as connectivity between cortical and subcortical circuitry matures (Casey et al., 2019; Parr et al., 2021). However, more neuroimaging work is needed to unveil how changes in neural circuitry support the use of these choice features during decision making across development.

Recent theories have proposed that heightened exploration is an adaptive quality of childhood (Giron et al., 2022; Gopnik, 2020), but these theoretical accounts — and the empirical work they have inspired — have left open the question of how different features of the environment elicit this strong, exploratory drive early in life. Here, we demonstrated that the developmental shift from more exploratory to more exploitative behavior may arise from strengthening aversion to reward uncertainty with increasing age, rather than from heightened sensitivity to novelty. Importantly, in the real world, exploration manifests through interaction with dynamic ecological contexts in which novelty and uncertainty may be highly idiosyncratic. Compared to adolescents and adults, children may encounter a greater number of novel options due to their relative lack of life experience. As they gain more autonomy, adolescents may find themselves facing more decisions (particularly in the social domain) with unknown reward outcomes. The current findings help build toward a comprehensive understanding of developmental change in exploration by disentangling the cognitive processes that govern how individuals interact with these core features of the choice options present in natural environments.

Data and Code Availability

Task code, anonymized data, and analysis code are available on the Open Science Framework: https://osf.io/cwf2k/

Acknowledgments

We thank May Levin for help with task design.

Funding

This work was supported by a National Science Foundation CAREER Award (Grant No. 1654393 to C.A.H.), the NYU Vulnerable Brain Project (grant to C.A.H.), the Department of Defense (NDSEG fellowship to K.N.), the National Institute of Mental Health (F31 MH129105 to K.N.; T32 MH019524 to R.E.M.), the National Institute of Child Health and Human Development (F31 HD097873 to R.E.M.), the Leon Levy Fellowship in Neuroscience (to R.E.M.), and the National Institute on Drug Abuse (F32 DA047047 to G.M.R.).

References

Badre, D., Doll, B. B., Long, N. M., & Frank, M. J. (2012). Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron*, *73*(3), 595–607. https://doi.org/10.1016/j.neuron.2011.12.025

Blanchard, T. C., & Gershman, S. J. (2018). Pure correlates of exploration and exploitation in the human brain. *Cognitive, Affective & Behavioral Neuroscience*, *18*(1), 117–126. https://doi.org/10.3758/s13415-017-0556-2

Blanco, N. J., & Sloutsky, V. M. (2020). Systematic exploration and uncertainty dominate young children's choices. *Developmental Science*, e13026. https://doi.org/10.1111/desc.13026

Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, *64*(4), 215–234. https://doi.org/10.1016/j.cogpsych.2011.12.002

Camerer, C., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. In *Journal of Risk and Uncertainty, 5*(4), 325–370. https://doi.org/10.1007/bf00122575

Casey, B. J., Heller, A. S., Gee, D. G., & Cohen, A. O. (2019). Development of the emotional brain. *Neuroscience Letters*, *693*, 29–34. https://doi.org/10.1016/j.neulet.2017.11.055

Christakou, A., Gershman, S. J., Niv, Y., Simmons, A., Brammer, M., & Rubia, K. (2013). Neural and psychological maturation of decision-making in adolescence and young adulthood. *Journal of Cognitive Neuroscience*, *25*(11), 1807–1823. https://doi.org/10.1162/jocn_a_00447

Cockburn, J., Man, V., Cunningham, W., & O'Doherty, J. P. (2021). Novelty and uncertainty interact to regulate the balance between exploration and exploitation in the human brain. *bioRxiv*. https://doi.org/10.1101/2021.10.13.464279

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. Series B, Biological Sciences, *362*(1481), 933–942. https://doi.org/10.1098/rstb.2007.2098

Daffner, K. R., Mesulam, M. M., Scinto, L. F., Cohen, L. G., Kennedy, B. P., West, W. C., & Holcomb, P. J. (1998). Regulation of attention to novel stimuli by frontal lobes: an event-related potential study. *Neuroreport*, *9*(5), 787–791. https://doi.org/10.1097/00001756-199803300-00004

Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, *173*, 34–42. https://doi.org/10.1016/j.cognition.2017.12.014

Giron, A. P., Ciranka, S. K., Schulz, E., van den Bos, W., Ruggeri, A., Meder, B., & Wu, C. M. (2022). Developmental changes in learning resemble stochastic optimization. *PsyArXiv.* https://doi.org/10.31234/osf.io/9f4k3

Gopnik, A. (2020). Childhood as a solution to explore–exploit tensions. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *375*(1803), 20190502. https://doi.org/10.1098/rstb.2019.0502

Gopnik, A., O'Grady, S., Lucas, C. G., Griffiths, T. L., Wente, A., Bridgers, S., Aboody, R., Fung, H., & Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences of the United States of America*. https://doi.org/10.1073/pnas.1700811114

Gottlieb, J., Oudeyer, P.-Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in Cognitive Sciences*, *17*(11), 585–593. https://doi.org/10.1016/j.tics.2013.09.001

Habicht, J., Bowler, A., Moses-Payne, M. E., & Hauser, T. U. (2021). Children are full of optimism, but those rose-tinted glasses are fading – reduced learning from negative outcomes drives hyperoptimism in children. *Journal of Experimental Psychology: General.* https://doi.org/10.1037/xge0001138

Henderson, B., & Moore, S. G. (1980). Children's Responses to Objects Differing in Novelty in Relation to Level of Curiosity and Adult Behavior. *Child Development*, *51*(2), 457–465. https://doi.org/10.2307/1129279

Houillon, A., Lorenz, R. C., Boehmer, W., Rapp, M. A., Heinz, A., Gallinat, J., & Obermayer, K. (2013). The effect of novelty on reinforcement learning. *Progress in Brain Research*, *202*, 415–439. https://doi.org/10.1016/B978-0-444-62604-2.00021-6

Jaegle, A., Mehrpour, V., & Rust, N. (2019). Visual novelty, curiosity, and intrinsic reward in machine learning and the brain. *Current Opinion in Neurobiology*, *58*, 167–174. https://doi.org/10.1016/j.conb.2019.08.004

Jepma, M., Schaaf, J. V., Visser, I., & Huizenga, H. M. (2020). Uncertainty-driven regulation of learning and exploration in adolescents: A computational account. *PLoS Computational Biology*, *16*(9), e1008276. https://doi.org/10.1371/journal.pcbi.1008276

Kakade, S., & Dayan, P. (2002). Dopamine: generalization and bonuses. *Neural Networks: The Official Journal of the International Neural Network Society*, *15*(4–6), 549–559. https://www.ncbi.nlm.nih.gov/pubmed/12371511

Krebs, R. M., Schott, B. H., Schütze, H., & Düzel, E. (2009). The novelty exploration bonus and its attentional modulation. *Neuropsychologia*, *47*(11), 2272–2281. https://doi.org/10.1016/j.neuropsychologia.2009.01.015

Lee, J. K., Wendelken, C., Bunge, S. A., & Ghetti, S. (2016). A Time and Place for Everything: Developmental Differences in the Building Blocks of Episodic Memory. *Child Development*, *87*(1), 194–210. https://doi.org/10.1111/cdev.12447

Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural representation of subjective value under risk and ambiguity. *Journal of Neurophysiology*, *103*(2), 1036–1047. https://doi.org/10.1152/jn.00853.2009

Liquin, E. G., & Gopnik, A. (2022). Children are more exploratory and learn more than adults in an approach-avoid task. *Cognition*, *218*, 104940.

https://doi.org/10.1016/j.cognition.2021.104940

Lloyd, A., McKay, R., Sebastian, C. L., & Balsters, J. (2020). Are Adolescents More Optimal Decision-Makers in Novel Environments? Examining the Benefits of Heightened Exploration in a Patch Foraging Paradigm. *Developmental Science, 24*(4), https://doi.org/10.1111/desc.13075

Meder, B., Wu, C. M., Schulz, E., & Ruggeri, A. (2021). Development of directed and random exploration in children. *Developmental Science*, *24*(4), e13095. https://doi.org/10.1111/desc.13095

Mendel, G. (1965). Children's Preferences for Differing Degrees of Novelty. *Child Development*, *36*(2), 453–465. https://doi.org/10.2307/1126468

Nussenbaum, K., & Hartley, C. A. (2019). Reinforcement learning across development: What insights can we draw from a decade of research? *Developmental Cognitive Neuroscience*, *40*, 100733. https://doi.org/10.1016/j.dcn.2019.100733

Parr, A. C., Calabro, F., Larsen, B., Tervo-Clemmens, B., Elliot, S., Foran, W., Olafsson, V., & Luna, B. (2021). Dopamine-related striatal neurophysiology is associated with specialization of frontostriatal reward circuitry through adolescence. *Progress in Neurobiology*, *201*, 101997. https://doi.org/10.1016/j.pneurobio.2021.101997

Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O'Doherty, J. P. (2013). The neural representation of unexpected uncertainty during value-based decision making. *Neuron*, *79*(1), 191–201. https://doi.org/10.1016/j.neuron.2013.04.037

Philpot, R. M., & Wecker, L. (2008). Dependence of adolescent novelty-seeking behavior on response phenotype and effects of apparatus scaling. In *Behavioral Neuroscience* (Vol. 122, Issue 4, pp. 861–875). https://doi.org/10.1037/0735-7044.122.4.861

Piray, P., Dezfouli, A., Heskes, T., Frank, M. J., & Daw, N. D. (2019). Hierarchical Bayesian inference for concurrent model fitting and comparison for group studies. *PLoS Computational Biology*, *15*(6), e1007043. https://doi.org/10.1371/journal.pcbi.1007043

Rich, A. S., & Gureckis, T. M. (2018). Exploratory choice reflects the future value of information. *Decisions*, *5*(3), 177–192. https://doi.org/10.1037/dec0000074

Rosenbaum, G. M., & Hartley, C. A. (2019). Developmental perspectives on risky and impulsive choice. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *374*(1766), 20180133. https://doi.org/10.1098/rstb.2018.0133

Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for Rewards Like a Child Means Less Generalization and More Directed Exploration. *Psychological Science*, *30*(11), 1561–1572. https://doi.org/10.1177/0956797619863663

Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, *43*(4), 1045–1050. https://doi.org/10.1037/0012-1649.43.4.1045

Shing, Y. L., Werkle-Bergner, M., Brehmer, Y., Müller, V., Li, S.-C., & Lindenberger, U. (2010). Episodic memory across the lifespan: the contributions of associative and strategic components. *Neuroscience and Biobehavioral Reviews*, *34*(7), 1080–1091. https://doi.org/10.1016/j.neubiorev.2009.11.002

Shing, Y. L., Werkle-Bergner, M., Li, S.-C., & Lindenberger, U. (2008). Associative and strategic components of episodic memory: a life-span dissociation. *Journal of Experimental Psychology. General*, *137*(3), 495–513. https://doi.org/10.1037/0096-3445.137.3.495

Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017). Charting the expansion of strategic exploratory behavior during adolescence. *Journal of Experimental Psychology. General*, *146*(2), 155–164. https://doi.org/10.1037/xge0000250

Spear, L. P. (2000). The adolescent brain and age-related behavioral manifestations. *Neuroscience and Biobehavioral Reviews*, *24*(4), 417–463. https://doi.org/10.1016/s0149-7634(00)00014-2

Stansfield, K. H., & Kirstein, C. L. (2006). Effects of novelty on behavior in the adolescent and adult rat. *Developmental Psychobiology*, *48*(1), 10–15. https://doi.org/10.1002/dev.20127

Sumner, E., Li, A. X., Perfors, A., Hayes, B., Navarro, D., & Sarnecka, B. W. (2019). *The Exploration Advantage: Children's instinct to explore allows them to find information that adults miss*. https://doi.org/10.31234/osf.io/h437v

Sutton, R. S., Barto, A. G., Co-Director Autonomous Learning Laboratory Andrew G Barto, & Bach, F. (1998). *Reinforcement Learning: An Introduction*. MIT Press. https://play.google.com/store/books/details?id=CAFR6IBF4xYC

Trudel, N., Scholl, J., Klein-Flügge, M. C., Fouragnan, E., Tankelevitch, L., Wittmann, M. K., & Rushworth, M. F. S. (2021). Polarity of uncertainty representation during exploration and exploitation in ventromedial prefrontal cortex. *Nature Human Behaviour*, *5*(1), 83–98. https://doi.org/10.1038/s41562-020-0929-3

Valenti, S. S. (1985). Children's preference for novelty in selective learning: Developmental stability or change? *Journal of Experimental Child Psychology*, *40*(3), 406–419. https://doi.org/10.1016/0022-0965(85)90074-8

Wang, J. (jenny), Yang, Y., Macias, C., & Bonawitz, E. (2021). Children With More Uncertainty in Their Intuitive Theories Seek Domain-Relevant Information. In *Psychological Science* (Vol. 32, Issue 7, pp. 1147–1156). https://doi.org/10.1177/0956797621994230

Wechsler, D. (2011). *Wechsler Abbreviated Scale of Intelligence (WASI-II)*. NCS Pearson.

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore--exploit dilemma. *Journal of Experimental Psychology. General*, *143*(6), 2074. https://doi/org/ 10.1037/a0038199

Wittmann, B. C., Daw, N. D., Seymour, B., & Dolan, R. J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron*, *58*(6), 967–973. https://doi.org/10.1016/j.neuron.2008.04.027