Exploration and Attention in Young Children

Nathaniel J. Blanco (nathanblanco@gmail.com)
Department of Psychology, The Ohio State University
1835 Neil Avenue, Columbus, OH 43210

Vladimir M. Sloutsky (Sloutsky.1@osu.edu)
Department of Psychology, The Ohio State University
1835 Neil Avenue, Columbus, OH 43210

Abstract

Exploration is critical for discovering how the world works. Exploration should be particularly valuable for young children, who have little knowledge about the world. Theories of decision-making describe systematic exploration as being primarily sub-served by prefrontal cortex (PFC). Recent research suggests that systematic exploration predominates in young children's choices, despite immature PFC, suggesting that this systematic exploration may be driven by different mechanisms. We hypothesize that young children's tendency to distribute attention widely promotes broad information gathering, which in turn translates to exploratory choice behavior, and that interrupting distributed attention allocation through bottom up attentional capture would also disrupt systematic exploration. We test this hypothesis using a simple choice task in which saliency of the options was manipulated. Saliency disrupted systematic exploration. These results suggest that attentional mechanisms may drive systematic exploratory behavior, and may be part of a larger tendency toward broad information gathering in young children.

Keywords: cognitive development; exploration; decision-making; attention

Introduction

One crucial way in which children's cognition differs from adults' is how they allocate their attention. Adults are highly adept at controlling their attention, distributing it broadly or focusing selectively on a small subset of stimuli (e.g., Chong & Treisman, 2005). In contrast, young children tend to distribute their attention broadly (Deng & Sloutsky, 2015, 2016; Smith & Kemler, 1977). In situations in which only a small amount of the available information is currently relevant, adults will often selectively focus on that piece of information, and ignore almost everything else (Rehder & Hoffman, 2005; Blair, Watson & Meier, 2009). Children, on the other hand, will distribute their attention to everything, even information that is not relevant for their current task or goals (Plebanek & Sloutsky, 2017). This tendency is likely to stem from immaturities of executive attention resulting in difficulty attending selectively and filtering out irrelevant environmental stimuli. While these immaturities may be highly limiting for learning in academic settings, it is possible that such immaturities of executive attention can be adaptive. For example, distributing attention can result in superior performance of children over adults in situations when participants have to use information that was previously irrelevant (Plebanek & Sloutsky, 2017; Blanco & Sloutsky, under review).

Therefore, depending on the context, either selective or distributed attention could be advantageous. Selective attention is superior when one is confident that a small portion of the available information is sufficient to achieve their goals. Distributed attention is advantageous when there is more uncertainty about what is and is not important. For young children, who have much less experience and knowledge about how things in the world work, it may often be the best strategy. Given their greater uncertainty about what is relevant, distributing attention is a safer bet. Additionally, by facilitating broad information gathering, it serves to reduce that uncertainty and build up the rich general knowledge that adults rely on (which, in turn, allows for effective use of selective attention later in life). In other words, it seems that distributing attention in young children might sub-serve exploration. Recent research suggests that there is a tight link between attention and decision-making behavior (Gottlieb, 2012; Konovalov & Krajbich, 2016). Perhaps distributed attention also promotes wider distribution in selection of potential choices. By distributing attention early in life children may be sacrificing immediate performance in exchange for information that can be used later.

There are recent reports indicating that four-year-old's choices are, indeed, highly exploratory (Blanco & Sloutsky, under review). Interestingly, children's exploration also appeared to non-random. This is surprising because research on exploratory behavior makes an important distinction between systematic and undirected (or random) exploration strategies (Badre, Doll, Long, & Frank, 2012; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; Knox, Otto, Stone, & Love, 2012; Blanco, Love, Cooper, McGeary, Knopik, & Maddox, 2015; Somerville, Sasse, Garrad, Drysdale, Abi Akar, Insel, & Wilson, 2017), and converging evidence suggests a crucial role of prefrontal cortex in systematic exploration (Badre, Doll, Long, & Frank, 2012; Frank, Doll, Oas-Terpstra, & Moreno, 2009; Blanco et al., 2015; Otto, Knox, Markman, & Love, 2014). Because prefrontal cortex exhibits substantially protracted development (Sowell, Thompson, Leonard, Welcome, Kan, & Toga, 2004; Sowell, Thompson, Holmes, Jernigan, & Toga, 1999), current theories predicted that young children's exploration would be largely unsystematic (Somerville et al., 2017). Due to the immaturity of PFC, young children's systematic exploration is likely driven by different mechanisms than adults'. We hypothesize that children's exploratory behavior is instead tied intricately to their immature attention allocation. exploration.

The Current Study

The goal of the current study is to test this idea by systematically manipulating saliency of a cue linked to a reward. More specifically our hypothesis is that children's typical pattern of distributed attention promotes distributing choices in a way that enables systematic exploration. If disrupting this process through bottom-up capture of attention by salient stimuli results in a change in exploratory behavior, we would be able to infer that attention early in development drives exploratory behavior. In contrast, if attention is not a causal factor in exploratory behavior, manipulating attention should lead to little or no changes in exploratory behavior.

In the current study, we presented children with a simple reward learning task under three attentional conditions in order to examine the interplay of attention and systematic exploration. On each trial of the task they chose between four options that gave different amounts of reward. The conditions differed in terms of the perceptual saliency of stimuli marking the choice options. In the Baseline condition, all options were of approximately equal salience. In the two experimental conditions, three choice options were represented by bland, invariant stimuli, while one option was represented by a highly salient stimulus that changed on every trial. In the *Congruent* condition the salient option was mapped to the option that gave the highest reward. In the *Competition* condition the salient option was mapped to the lowest reward, putting reward-seeking and salience in competition with each other.

Methods

Participants A total of 110 four- and five-year-olds (mean age = 57 months; 58 girls) participated in the experiment: 37 in the congruent condition, 37 in the competition condition, and 36 in the baseline condition. Participants were recruited from preschools and childcare centers in the Columbus, Ohio area.

Procedure Participants completed a simple decision-making task that was framed as a computer game in which they asked alien creatures for candy (Figure 1). The task was a simplified version of a standard *n*-armed bandit task commonly used to study reward-based decision-making (e.g., Daw et al. 2006). The goal of the game was to earn as much candy as possible. On each of 100 trials, participants chose one out of the four creatures and received (virtual) candy for their choice. Selections were made using a touch screen. Each creature gave a set number of candies that was the same on every trial: One option was 10 candies, while the other three options were 3, 2, and 1 candies respectively. The locations of the reward values were stable across the entire experiment but were randomly determined for each participant. Following the choice, the reward received for the choice was displayed for 3 s (Figure 1B). Then a meter that tracked the total accumulated reward was updated. Children were given tangible rewards (stickers) for every 180 candies earned, with benchmarks on the meter indicating these goals. When a goal was reached, a congratulatory screen appeared telling the participant that they earned a sticker. It should be noted that, because outcomes were stable and predictable, low levels of exploration would usually be expected, but we have previously found high levels of systematic exploration in young children in this task (Blanco & Sloutsky, under review).

Participants were assigned to one of three conditions: Congruent, Competition, and Baseline, with salience manipulated across the conditions. In the Baseline condition, all creatures were approximately equally salient, whereas in the Congruent and Competition conditions, salience of the creatures was unequal. Specifically, three of the four creatures were simple black and white stick figures, whereas one was colorful and perceptually rich. In addition, on each trial the salient image was a different novel creature (Figure 1C). Fifty different images were used, so each image appeared twice during the experiment. In the Congruent condition, the salient option was mapped to the highest reward value (10 candies), whereas in the Competition condition, the salient option was mapped to the lowest reward value (1 candy).

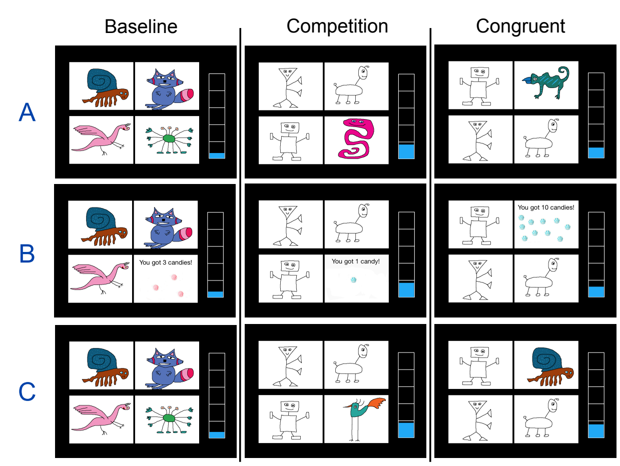

Figure 1: Trial structure. (A) After each choice, (B) the reward earned for the choice is presented for 3 s, (C) then the next trial begins. In the Congruent and Competition conditions one option is represented by a colorful image that changes on every trial, while the other three are represented by lower saliency images that remain stable across trials. In the Baseline condition, all four options are represented by stable images of equal saliency.

Results

Choice Proportions Participants' choices over the course of the experiment are presented in Figure 2. The analysis focused on the proportion of trials on which the highest valued option was chosen. The main purpose of this analysis was to assess the effect of saliency on performance across the three conditions. An ANOVA revealed a significant effect of condition, $F(2, 107) = 15.40$, $p = 0.001$, $\eta^2 = 0.22$. Pairwise comparisons showed that participants in the Congruent condition ($M = 0.53$) chose the best option (i.e., the 10-candy option) significantly more often than participants in the Baseline condition ($M = 0.28$), $t(71) = 4.40$, $p < 0.001$, $d = 1.03$, and the Competition condition ($M = 0.30$), $t(72) = 3.93$, $p < 0.001$, $d = 0.91$. Performance in the Competition condition was not different than in the Baseline condition, $t(71) = 0.57$, $p = 0.569$, $d = 0.13$.

The proportion of trials in which the lowest valued option was chosen was also analyzed in order to assess the effect of saliency in the Competition condition. An ANOVA revealed a significant effect of condition, $F(2, 107) = 5.24$, $p = 0.006$, $\eta^2 = 0.09$. But, pairwise tests revealed only that participants in the Congruent condition ($M = 0.16$) chose the lowest option less than both the Baseline ($M = 0.23$), $t(71) = 3.22$, $p = 0.002$, $d = 0.75$, and the Competition condition ($M = 0.25$), $t(72) = 2.60$, $p = 0.011$, $d = 0.60$. The Competition condition and the Baseline condition did not differ significantly, $t(71) = 0.65$, $p = 0.516$, $d = 0.15$. This pattern of results suggests that perceptual saliency facilitated reward optimization in the Congruent condition, but not through simple novelty-seeking since the salient option was not selected more frequently than Baseline in the Competition condition.

Switch Proportions We also examined the proportion of trials on which participants switched responses, choosing a different option than on the previous trial (see Figure 3). In the Baseline we expected participants to switch often and do so systematically. This expectation is based on a previous study (Blanco & Sloutsky, under review) using a task similar to the Baseline condition. In that study, children tended to switch extremely often—consistent with highly elevated exploration levels. Systematicity in their switching was then established with subsequent computational modeling. We, therefore, first analyze participants' behavioral responses. In the next section, we report modeling results.

An ANOVA on proportion of trials that participants switched responses revealed a significant effect of condition, $F(2, 107) = 17.42$, $p < 0.001$, $\eta^2 = 0.246$. Most importantly, Children in the Congruent condition ($M = 0.56$), $t(71) = 5.57$, $p < 0.001$, $d = 1.30$ and in the Competition condition ($M = 0.77$), $t(71) = 3.22$, $p = 0.002$, $d = 0.75$ exhibited substantially less switching than in the Baseline condition ($M = 0.91$). Additionally, children in the Competition condition switched more than those in the Congruent condition, $t(72) = 3.04$, $p = 0.003$, $d = 0.71$. It is perhaps not surprising that children switched less in the Congruent condition compared to Baseline since they are often exploiting the best option, but it is surprising that switching is low in the Competition condition despite no increase in exploitation.

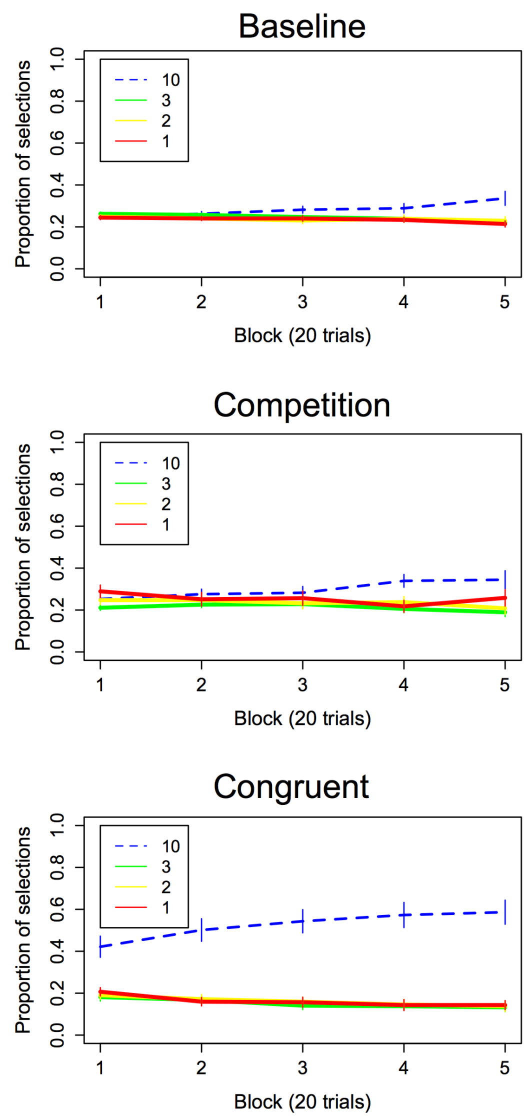

Figure 2: Choice proportions. The proportion of trials on which each option was chosen is presented for blocks of 20 trials. Compared to baseline, children in the congruent condition selected the highest valued option more frequently. Children in the competition condition selected the highest valued option less often that either the baseline or congruent conditions. Interestingly, children in the competition condition did not select the lowest valued option (which was salient in that condition) more often than in the other conditions where it was less salient. This suggests that pure novelty/saliency seeking did not drive children's choices. Error bars reflect standard errors of the mean.

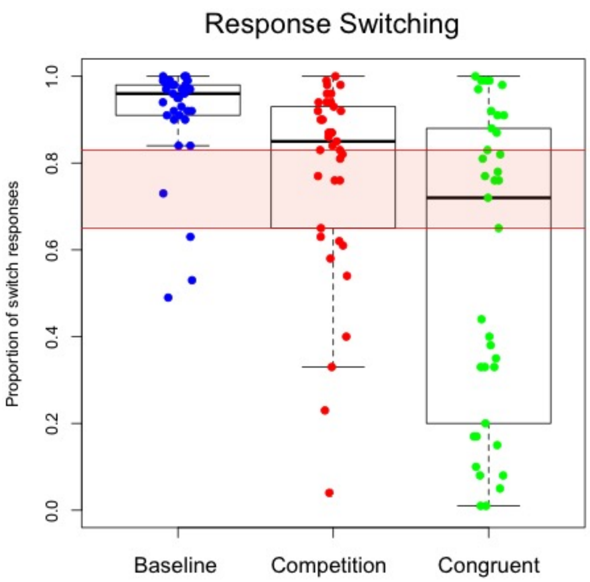

Figure 3: Response Switching. The proportion of trials on which participants made a switch response, choosing a different option than the previous trial, is presented. The pink shaded region represents 95% probability density of switch responses given random responding. Extreme switch proportions in the Baseline condition suggest elevated exploration levels. Switch proportions are less than the Baseline in both Saliency conditions. Dots represent individual participants.

Computational Modeling

In order to better understand children's choice strategies, and to examine the effect of the saliency manipulation of directed exploration, participants' choices were evaluated in relation to a Reinforcement Learning model (Sutton & Barto, 1998) that included the potential for both systematic (or directed) and random exploration. The model learned the reward values by updating expected values for each option based on the prediction error using the following equation:

$$V_{i,t+1} = V_{i,t} + \alpha\left(R_{i,t} - V_{i,t}\right)$$

where $V_{i,t}$ is the expected value of option i on trial t, $R_{i,t}$ is the reward is the reward on trial t earned for choosing option i, and α is the learning rate (a free parameter). It then made choices according to the following function:

$$P\left(a_{i,t}\right) = \frac{e^{\beta*[V_{i,t}*(1-\phi)+L_{i,t}*\phi]}}{\sum_{j=1}^{n} e^{\beta*[V_{j,t}*(1-\phi)+L_{j,t}*\phi]}}$$

where $P(a_{i,t})$ is the probability of choosing option i on trial t. $L_{i,t}$ is the lag term—a proxy for uncertainty—that simply encodes the number of trials since option i was last chosen. ϕ is the weight parameter mediating the relative extent to which the expected values and lags influence choices and is constrained to be between 0 and 1, inclusive. Greater values of ϕ indicate greater influence of systematic exploration. When ϕ is 0, the model chooses based only on expected value; when ϕ is 1 it chooses only based on the lag. β is the inverse temperature parameter that controls random exploration. At $\beta = 0$ choice probabilities become completely random (i.e. equal between all options irrespective of value or lag), and as β approaches infinity the model chooses the most favorable option (based on the weighted combination of expected value and lag described above) on every trial. Both β and ϕ were free parameters.

This model is similar to the 'exploration bonus' models used in some previous studies (Daw et al. 2006; Kakade & Dayan, 2002), but with lag as a proxy for uncertainty and with slightly different parameterization. The model was fit to each individual participant by finding the set of parameters that maximized the likelihood of producing the participant's data given the model.

The full model described above was first compared to a simplified model that did not include systematic exploration, where ϕ was set equal to 0. In the simplified model choice probabilities reduce to a standard Softmax choice rule (Sutton & Barto, 1998) on expected reward value:

$$P\left(a_{i,t}\right) = \frac{e^{\beta*V_{i,t}}}{\sum_{j=1}^{n} e^{\beta*V_{j,t}}}$$

The Aikaike Information Criterion (AIC) was used to determine best-fitting model for each participant (Akaike, 1974). A large majority of children in the Baseline (31 out of 36) condition were better fit by the full model that included systematic exploration, while only about half of children were better fit by the full model in both the Congruent (18 out of 37) and Competition (20 out of 37) conditions. A chi-squared test confirmed that these proportions were different, X^2 (2; N = 110) = 12.75, p = 0.002.

The best-fitting parameter values were also compared between the three different conditions (see Figure 4 for distribution of parameter values across the conditions). Because data in Figure 4 suggest that the best-fitting parameter values were not normally distributed, we report median values (see Table 1) and compare groups using Wilcoxon rank sum tests. Best-fitting ϕ parameter (reflecting systematic exploration) was significantly lower in the Congruent, W = 1044, p < 0.001 and the Competition conditions, W = 886, p = 0.015 than in the Baseline condition. The Congruent and Competition conditions did not differ in best-fitting ϕ value, W = 606, p = 0.401. The β parameter was lower in the Competition condition compared to both the Baseline, W = 935, p = 0.002, and the Congruent conditions, W = 906, p = 0.016. The Baseline and Congruent conditions were not different, W = 743, p = 0.400.

The high values of ϕ in the Baseline condition suggest a large influence of systematic exploration on participants' choices, while the substantially lower values of ϕ in both the Congruent and Competition conditions suggest much lower levels of systematic exploration in these two conditions. While the choice proportions suggest that in the Congruent condition systematic exploration was largely replaced with reward maximizing choices (i.e. exploitation), in the Competition condition the low value of β suggests a greater amount of random exploration in that condition. Importantly, the salience manipulation dramatically decreased systematic exploration in the Congruent and Competition conditions compared to the Baseline, although this reduction was achieved by different mechanism – through increased exploitation in the Congruent condition and through increased random explorations in the Competition condition. These results suggest a direct link between attention and exploratory behavior early in development.

Table 1: Median best-fitting parameter values (with standard deviations in parentheses)

	Baseline	Competition	Congruent
ϕ	0.701	0.213 (0.40)	0.045 (0.31)
β	1.348 (6.1)	0.367 (12)	0.680 (876)*

*Note: While the medians suggest most values were small, β has no upper limit, and infrequent large outliers (which are consistent with reward maximization) result in large standard deviations. This large value is mainly due to two such outliers, without which the standard deviation of the remaining sample is 12.

Discussion

In this study, we examined the effects of an attentional manipulation on young children's choices and exploratory behavior. The goal of the current study was to examine the link between systematic exploration and attention by manipulating attention and observing effects of such manipulation on systematic exploration. The reported results suggest that attentional manipulation (i.e., exogenously capturing attention through large differences in salience) decreased the level of systematic exploration that young children exhibited in the Baseline condition.

In addition, children's choices did not indicate that they were simply salience-seeking in their choice strategy. Instead the effect of saliency was dependent on whether or not it was congruent with or in competition with reward maximization. When the salient option was also highly valuable, children chose that option more often than in the other conditions. But, when the salient option was low in value, it was not chosen any more often than in the other conditions. This interaction suggests a complex role of attention in determining young children's choices.

Together these results suggest that attentional mechanisms are a major determinant of exploratory behavior in young children. When saliency and/or novelty are otherwise equal, systematic exploration dominates choices (as in the Baseline condition), with options that were less recently sampled being more likely to be selected. A manipulation affecting bottom-up attentional seems to disrupt this process, leading to a reduction in systematic exploration. Once disrupted, if the salient stimulus signals a rewarding option, it can act as a cue that facilitates reward learning.

These results point to an integral role of attentional mechanisms in systematic exploratory behavior in young children, in contrast to the top-down PFC mediated processes involved in systematic exploration in adults. Despite PFC being immature, young children's tendency to distribute attention seems to support systematic exploration. Attentional mechanisms and exploratory decision-making may be part of a larger general pattern in young children in which their cognition and behavior are specifically tuned to facilitate broad information gathering—something that is particularly critical early in life.

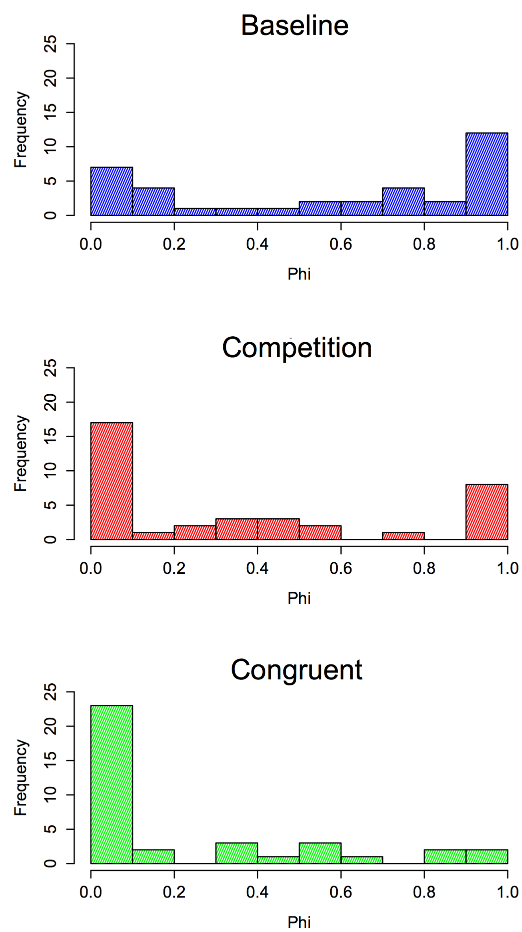

Figure 4: Best-fitting ϕ parameter. Histograms of the best-fitting ϕ parameter for each group are presented. Both saliency conditions have a large proportion of participants with very low values of ϕ, indicating little systematic exploration, while the Baseline condition has a larger proportion of participants with high values of ϕ, indicating greater systematic exploration.

Acknowledgments

This research was supported by National Institutes of Health Grants R01HD078545 and P01HD080679 to V. M. Sloutsky.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE transactions on automatic control*, *19*(6), 716-723.

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.

Blair, M. R., Watson, M. R., & Meier, K. M. (2009). Errors, efficiency, and the interplay between attention and category learning. *Cognition, 112*(2), 330-336.

Blanco, N. J., Love, B. C., Cooper, J. A., McGeary, J. E., Knopik, V. S., & Maddox, W. T. (2015). A frontal dopamine system for reflective exploratory behavior. *Neurobiology of learning and memory, 123*, 84-91.

Blanco & Sloutsky (under review). Advantages of developmental limitations: Exploration dominates choices in young children.

Blanco & Sloutsky (under review). Adaptive Flexibility in Category Learning? Young Children Exhibit Smaller Costs of Selective Attention than Adults.

Chong, S. C., & Treisman, A. (2005). Attentional spread in the statistical processing of visual displays. *Perception & Psychophysics*, *67*, 1–13.

Daw, N. D., O'doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature, 441*(7095), 876-879.

Deng, W. S., & Sloutsky, V. M. (2015). The development of categorization: Effects of classification and inference training on category representation. *Developmental Psychology, 51*(3), 392-405.

Deng, W. S., & Sloutsky, V. M. (2016). Selective attention, diffused attention, and the development of categorization. *Cognitive Psychology, 91,* 24-62.

Frank, M.J., Doll, B.B., Oas-Terpstra, J., & Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature Neuroscience, 12*, 1062–1068.

Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron, 76*(2), 281-295.

Kakade, S., & Dayan, P. (2002). Dopamine: Generalization and bonuses. *Neural Networks, 15,* 549–559.

Konovalov, A., & Krajbich, I. (2016a). Gaze data reveal distinct choice processes underlying model-based and model-free reinforcement learning. *Nature communications, 7*.

Knox, W. B., Otto, A. R., Stone, P., & Love, B. (2012). The nature of belief-directed exploratory choice in human decision-making. *Frontiers in psychology, 2,* 398.

Otto, A. R., Knox, W. B., Markman, A. B., & Love, B. C. (2014). Physiological and behavioral signatures of reflective exploratory choice. *Cognitive, Affective, & Behavioral Neuroscience, 14,* 1167–1183.

Plebanek, D. J., & Sloutsky, V. M. (2017). Costs of selective attention: when children notice what adults miss. *Psychological Science, 28*(6), 723-732.

Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology, 51*(1), 1-41.

Smith, L. B., & Kemler, D. G. (1977). Developmental trends in free classification: Evidence for a new conceptualization of perceptual development. *Journal of Experimental Child Psychology, 24*(2), 279-298.

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.

Sowell, E. R., Thompson, P. M., Holmes, C. J., Jernigan, T. L., & Toga, A. W. (1999). In vivo evidence for post-adolescent brain maturation in frontal and striatal regions. *Nature neuroscience, 2*(10), 859-861.

Sowell, E. R., Thompson, P. M., Leonard, C. M., Welcome, S. E., Kan, E., & Toga, A. W. (2004). Longitudinal mapping of cortical thickness and brain growth in normal children. *Journal of Neuroscience, 24*(38), 8223-8231.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction.* Cambridge: MIT press.