

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Developmental Change in What Elicits Curiosity

Permalink

<https://escholarship.org/uc/item/43g7m167>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

Authors

Liquin, Emily G

Callaway, Frederick

Lombrozo, Tania

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Developmental Change in What Elicits Curiosity

Emily G. Liquin (eliquin@princeton.edu)
Frederick Callaway (fredcallaway@princeton.edu)
Tania Lombrozo (lombrozo@princeton.edu)

Department of Psychology, Princeton University
Princeton, NJ 08540 USA

Abstract

Across the lifespan, humans direct their learning towards information they are curious to know. However, it is unclear what elicits curiosity, and whether and how this changes across development. Is curiosity triggered by surprise and uncertainty, as prior research suggests, or by expected learning, which is often confounded with these features? In the present research, we use a Bayesian reinforcement learning model to quantify and disentangle surprise, uncertainty, and expected learning. We use the resulting model-estimated features to predict curiosity ratings from 5- to 9-year-olds and adults in an augmented multi-armed bandit task. Like adults' curiosity, children's curiosity was best predicted by expected learning. However, after accounting for expected learning, children (but not adults) were also more curious when uncertainty was higher and surprise lower. This research points to developmental changes in what elicits curiosity and calls for a reexamination of research that confounds these elicitors.

Keywords: curiosity; learning; development; exploration

Children come into the world with a great deal to learn, and they actively direct their learning by selectively attending, exploring, and asking questions. Often, such information search is motivated by *curiosity*, the phenomenological experience of wanting to know (Liquin & Lombrozo, 2020b). Curiosity not only directs learning, but can potentially improve it: older children and adults better remember facts about which they were more curious (Fandakova & Gruber, 2021; Kang et al., 2009), and young children learn object labels better after pointing at the object (perhaps indicating curiosity; Lucca & Wilbourn, 2018). While targets of curiosity may enjoy an advantage for learning, it is not clear how those targets are chosen—especially by young children. In the current paper we ask: what elicits curiosity, and how might this change over the course of development?

Though curiosity might enhance learning when new and valuable information is obtained, many quests for information fail: we do not always find new and useful information when we seek it. In order to maximize the chance of successfully achieving learning, an ideal learner should experience curiosity about—and hence direct their information search towards—queries that have the greatest potential to yield such information. In the present research, we test whether children and adults are “optimal” in this sense, experiencing curiosity selectively when learning is likely to occur. We contrast this optimal strategy with a simpler strategy: learners might instead use indirect “heuristic cues” to

expected learning, such as surprise and uncertainty—two plausible elicitors of curiosity based on prior research.

Surprise and uncertainty are “heuristic cues” because they are reasonable but imperfect guides to expected learning, and because they are likely easier to gauge. Surprise and uncertainty often co-occur with expected learning: for example, a magic trick might elicit surprise, uncertainty about how the trick worked, and the expectation that one could learn by questioning the magician. However, surprise and uncertainty can also lead a learner astray. For example, a child might watch cars driving past and wonder whether their colors form a pattern. As each new car appears, the child will be surprised by its color and uncertain about the pattern—but there is no pattern to be found, and thus no opportunity for learning. Moreover, surprise and uncertainty may be easier to compute than expected learning: the former rely solely on prior beliefs, while the latter also incorporates beliefs about the future.

Do children and adults experience curiosity when learning is likely to occur, or do they rely on heuristic cues (surprise and uncertainty)? Though these potential elicitors of curiosity might often co-occur, disentangling their influence on curiosity can shed light on the mechanisms that drive curiosity and offer the potential to develop interventions that elicit curiosity when it would be beneficial (e.g., in educational settings). Below we briefly review prior work on what elicits children's and adults' curiosity, which largely does not differentiate between the optimal strategy of tracking expected learning and the use of simpler heuristic cues. We then introduce a method that allows us to tease apart expected learning, surprise, and uncertainty. To preview our results, we find that adults' curiosity tracks expected learning almost exclusively, whereas children's curiosity tracks expected learning less strongly, with additional roles for surprise and uncertainty.

Triggers of Curiosity

Prior research has not typically measured curiosity directly, especially in children. Instead, prior work has investigated the cues that shape attention, exploration, and question asking, behaviors that are likely to be at least in part triggered by curiosity. This research has suggested an important role for surprise (violation of expectation) and uncertainty in eliciting information search. Children explore and attend to objects that violate their expectations, whether those expectations are based on core knowledge, intuitive theories, or probabilistic evidence (Bonawitz et al., 2012; Kidd et al., 2012; Sim & Xu, 2017; Stahl & Feigenson, 2015). Moreover, children

preferentially explore when they are uncertain—for example, when two equally likely hypotheses are compatible with observed evidence (Cook et al., 2011; Schulz & Bonawitz, 2007). In adults, surprise and uncertainty are associated with both self-report measures of curiosity (Kang et al., 2009; Vogl et al., 2020) and behavioral measures of information search (Itti & Baldi, 2009; Kobayashi et al., 2019).

For adults, there is also evidence that curiosity tracks expected useful learning (Abir et al., 2020; Dubey & Griffiths, 2020), above and beyond heuristic cues including surprise and uncertainty (Liquin et al., 2020; Liquin & Lombrozo, 2020a). In contrast, research on children’s information search has not disentangled surprise and uncertainty from expected learning. For example, a child who observes a surprising event (e.g., a ball rolling through a solid wall; Stahl & Feigenson, 2015) might also expect to learn by exploring. Because surprise and uncertainty are confounded with expected learning, it remains unclear what elicits children’s curiosity. Are children more curious when they expect to learn? Or does children’s curiosity instead track surprise and uncertainty, which may be reasonable (but sometimes unreliable) cues to when learning is likely to occur?

Two lines of research offer initial insight into these questions. First, recent studies have disentangled some heuristic cues to expected learning. In one study, infants did not explore after a surprising event that was immediately explained away (Perez & Feigenson, 2020), suggesting a potential role for uncertainty or expected learning above and beyond surprise. Another study (Poli et al., 2020) found that surprise, uncertainty, and learning each explained variance in infants’ attention—but the authors tested learning from the just-observed trial (i.e., previous learning) rather than learning from future exploration (i.e., expected learning). Like surprise and uncertainty, previous learning may be a cue to expected learning, but these quantities are distinct.

Other research has investigated how children select which question to ask (which could be partly determined by curiosity), focusing on expected learning as the optimal determinant of a question’s quality (for a review, see Jones et al., 2020). Like adults (Rothe et al., 2018), preschoolers can select the optimal question from a small set of alternatives (Ruggeri et al., 2017), suggesting that children might have the capacity to detect expected learning. However, this research has not contrasted expected learning with surprise and uncertainty.

The Present Research

In the present research, we investigate to what extent children’s curiosity tracks surprise, uncertainty, and expected learning. Building on an experimental paradigm and computational model introduced by Liquin et al. (2020), we operationalize these candidate triggers precisely, allowing us to tease apart their roles in guiding curiosity.

In our reinforcement learning task (adapted from Dorfman et al. 2019), participants choose between two options (candy machines) over 20 trials. Each option has a fixed probability of producing a rewarding outcome (candy). However, there is an additional causal influence, such that each outcome can

be caused by either the option itself or by an intervening agent. On each trial, the participant reports their curiosity about the cause of the outcome (the machine or the agent). Using a Bayesian reinforcement learning model, we quantify surprise, uncertainty, and expected learning on each trial, allowing us to test whether curiosity tracks each feature.

In addition to testing several potential triggers of children’s curiosity, we compare children’s curiosity to that of adults. Computing expected learning is a complex operation in terms of its cognitive and metacognitive demands. Perhaps reflecting these demands, 7- to 10-year-olds are more likely than adults to ask questions that do not provide any new information (Ruggeri et al., 2016) and are less attuned to expected learning in their exploration (Nussenbaum et al., 2020). Similarly, we might expect children’s curiosity to be less attuned to expected learning (and perhaps more attuned to surprise and uncertainty) than adults’ curiosity. To test this, we compare adults with 5- to 9-year-olds, an age range in which children show some sophistication in question-asking ability yet still differ from adults (Jones et al., 2020).

We measure curiosity through self-report. As we previously suggested, decisions to direct one’s attention, explore, or ask a question might be partly motivated by curiosity. However, these behaviors can arise even when curiosity is not present (e.g., a child might ask a question to facilitate a social exchange), and curiosity can be experienced but not pursued (e.g., a child might wonder about an object’s label but not ask their parent about it). As a result, the study of information search does not provide conclusive evidence regarding the elicitors of curiosity. Thus, departing from most prior research, we attempt to measure curiosity more directly, by asking participants to give explicit ratings of how curious they are (i.e., how much they “want to know”).

Computational Model

We modeled the hypothesized determinants of curiosity using a Bayesian reinforcement learning model developed by Liquin et al. (2020). The model tracks the probability of receiving candy or no candy from a given candy machine.

On each trial, t , the candy machine generates candy with probability θ . However, a hidden agent intervenes with probability ε (with $\varepsilon = 1/3$). The intervention is modeled as a latent variable, $Z_t \sim \text{Bernoulli}(\varepsilon)$, where $Z_t = 1$ means an intervention occurred on trial t . When an intervention does occur, candy is instead produced with probability θ_z , which is set to 1 or 0 depending on the experimental condition. As a result, the reward R_t for a trial is distributed: $R_t \sim \text{Bernoulli}(\theta(1 - Z_t) + \theta_z Z_t)$.

The model estimates θ , that is, the probability of a machine producing candy. We assume a uniform prior on θ between 0 and 1. The posterior is computed by marginalizing over the sequence of interventions, $\vec{z}$,

$$p(\theta \mid r_{1:t}, \theta_z, \varepsilon) \propto \sum_{\vec{z}} \prod_t p(z_t \mid \varepsilon) p(r_t \mid \theta, \theta_z, z_t).$$

As noted by Liquin et al. (2020), marginalizing over $\vec{z}$ directly is intractable, but the posterior can nonetheless be computed because it depends only on the number of times each

Table 1: Definition of each model estimated feature, and associations with curiosity in simple linear regression models.

Feature	Definition	β [95% CI]	
		Children	Adults
Surprise-IT	Unlikelihood of received reward, based on machine's estimated value and possibility of intervention	-0.21 [-0.27, -0.16]	-0.28 [-0.34, -0.23]
Surprise-Full	Difference between received and expected reward based on estimated value and possibility of intervention	-0.20 [-0.26, -0.14]	-0.28 [-0.34, -0.23]
Surprise-Mean	Difference between received and expected reward based on mean estimate of the machine's value	-0.005 [-0.06, 0.05]	0.20 [0.14, 0.25]
Surprise-MAP	Difference between received and expected reward based on most likely estimate of the machine's value	-0.08 [-0.13, -0.02]	0.12 [0.06, 0.18]
Value Uncertainty	Uncertainty about the machine's value	0.08 [0.02, 0.14]	0.11 [0.05, 0.16]
Query EIG/Uncertainty	Expected learning/uncertainty about whether an intervention occurred (the target query)	0.27 [0.21, 0.32]	0.53 [0.48, 0.57]
Value EIG	Expected learning about chosen machine's value if the target query were to be answered	0.25 [0.19, 0.31]	0.52 [0.47, 0.57]

possible combination of (r, z) occurs. We assume that θ_z and ε are known. For concision, we omit the dependence on these parameters in later equations.

This model estimates θ for a given candy machine; the model is copied to estimate θ for a second candy machine. On each trial, in addition to estimating θ , the model computes a number of quantities, defined below (see also Table 1).

Surprise features capture the extent to which the outcome on a given trial is unexpected in light of prior beliefs. We consider four possible definitions of surprise. One definition captures information-theoretic surprisal (Shannon, 1948) which has been previously used to model infants' attention (Kidd et al., 2012; Poli et al., 2020). The other three definitions capture unsigned reward prediction error, which has been previously used to model adults' attention (Stojić et al., 2020). Because both information-theoretic surprisal and unsigned reward prediction error have been used in prior research, we model both.

Information-theoretic surprisal (*surprise-IT*) is defined as the negative log posterior predictive probability of the observed reward given the history of previous rewards, $-\log p(r_t | r_{1:t-1})$. The posterior predictive is defined

$$p(r_t | r_{1:t-1}) = \int_0^1 p(\theta | r_{1:t-1}) \sum_{z_t} p(z_t) p(r_t | \theta, z_t) d\theta.$$

Unsigned reward prediction error is defined $|R_t - \bar{r}_t|$, or the absolute difference between expected reward and received reward. We consider three possible definitions of expected reward $\bar{r}_t$. First, the full Bayesian estimate, taking into account both the current estimate of θ and the possibility of intervention, is $\bar{r}_t^{\text{full}} = p(R_t = 1 | r_{1:t-1})$. Two additional definitions do not consider the possibility of intervention: the posterior mean estimate of θ prior to the current reward, $\bar{r}_t^{\text{mean}} = \int_0^1 p(\theta | r_{1:t-1}) \theta d\theta$ and the maximum a posteriori (MAP) estimate of θ prior to the current reward, $\bar{r}_t^{\text{MAP}} = \text{argmax}_{\theta} p(\theta | r_{1:t-1})$. This results in three additional

surprise features: *surprise-full*, *surprise-mean*, and *surprise-MAP*.

Uncertainty features describe uncertainty about model estimates, defined using the information-theoretic measure entropy (Shannon, 1948). The first feature, *value uncertainty*, captures uncertainty about the estimated value of the chosen machine (i.e., its probability of producing candy, θ). This is defined $H(\theta | r_{1:t}) = -\int_0^1 p(\theta | r_{1:t}) \log(p(\theta | r_{1:t})) d\theta$. The second feature, *query uncertainty*, captures uncertainty about whether an intervention occurred (i.e., the target query about which participants rate their curiosity). This is defined as the entropy of the intervention's predictive distribution, $H(Z_t | r_{1:t}) = -\bar{z}_t \log \bar{z}_t - (1 - \bar{z}_t) \log(1 - \bar{z}_t)$ where $\bar{z}_t = p(Z_t = 1 | r_{1:t})$ is the conditional probability of the agent having intervened on this trial given all observed rewards,

$$\bar{z}_t = \int_0^1 p(\theta | r_{1:t}) \frac{p(r_t | \theta, Z_t = 1) p(Z_t = 1)}{\sum_{z_t \in \{0,1\}} p(r_t | \theta, z_t) p(z_t)} d\theta.$$

Expected learning features encode how much learning is expected to occur if the target of one's curiosity (whether an intervention occurred) were to be revealed. We define expected learning using a measure of expected information gain (EIG) that has been influential in the study of inquiry (see Coenen et al., 2019). First, expected learning about the chosen machine's value, or *value EIG*, is the expected reduction in entropy of the posterior distribution of θ after observing the value of Z_t , $\sum_{z \in \{0,1\}} p(Z_t = z | r_{1:t}) H(\theta | r_{1:t}, Z_t = z) - H(\theta | r_{1:t})$. Second, expected learning about whether an intervention occurred, or *query EIG*, is the expected reduction in entropy in the predictive distribution, $H(Z_t | r_{1:t})$, after observing the value of Z_t . Because observing the value of Z_t would reduce entropy to zero, query EIG is equal to query uncertainty, and we refer to this as "query EIG/uncertainty."

To summarize, though surprise, uncertainty, and expected learning often co-occur (see Fig. 1), the provided mathematical definitions allow us to pull apart these features, including several possible definitions of surprise and several possible

targets of uncertainty/expected learning. In the following section, we describe an experiment in which we test how well these features predict children’s and adults’ curiosity.

Experiment

Method

Participants We recruited 55 five- to nine-year-old children (2 five-year-olds, 13 six-year-olds, 14 seven-year-olds, 14 eight-year-olds, and 12 nine-year-olds; 32 female, 23 male), using a laboratory registry and word-of-mouth. We also recruited 55 adults (ages 18-62; 33 female, 20 male, one non-binary, and one unspecified) from Prolific, who were required to be in the United States and have a history of participating in at least 100 studies with a minimum 95% approval rate. An additional 12 children and five adults were excluded for providing the same curiosity rating on all trials.

Procedure All children and adults completed the experiment on a computer. Child participants were supervised by their parents. The full procedure was narrated by a pre-recorded “experimenter” and illustrated with animated videos.

The procedure was modified from Dorfman et al. (2019) and Liquin et al. (2020). Participants were introduced to virtual candy machines, which sometimes produced virtual candy when clicked. Some machines “work really well” and provided candy often, while other machines “don’t work so well.” Participants practiced clicking on two machines, which differed in their probability of producing candy.

Participants were then introduced to a squirrel named AJ, who interferes with the candy machines’ performance. In the *donor condition* ($N = 25$ children, 30 adults), participants were told that AJ likes to give people candy, while in the *thief condition* ($N = 30$ children, 25 adults), participants were told that AJ likes to take away people’s candy. Participants learned through practice with two machines that AJ intervened on one-third of trials; when AJ intervened, they received candy (in the donor condition) or no candy (in the thief condition) with 100% probability.

During this training, participants were also introduced to a curiosity rating scale. On the first practice trial after learning about AJ, the experimenter narrated, “Huh, I wonder whether we [got candy/didn’t get any candy] because AJ [put it there/took it away], or because of how well the machine works!” (with the text in brackets dependent upon condition). The experimenter then introduced a four-point rating scale, based on the idea that “in school, you raise your hand when you want to know something.” Participants were shown four cartoon images of raised hands, varying in size. For each hand, the experimenter described varying levels of “wanting to know whether AJ made that happen,” with higher levels of “wanting to know” corresponding to larger hands. Participants practiced selecting a hand to indicate their curiosity on multiple trials and answered several questions assessing their understanding of the rating scale (with corrective feedback).

Participants then advanced to the main task. Participants were shown two new candy machines, which differed in color




Figure 1: Pearson correlations between all model-estimated features, over all trials for both child and adult participants.

(blue or green) and were visually distinct from the practice machines. Participants were instructed that they needed to learn how well the machines worked over 20 trials. On each trial, participants chose to click on one machine, receiving either candy or no candy. After each outcome, participants rated their curiosity about whether AJ had intervened. Whether or not AJ intervened was not revealed. Across all trials, the machines produced candy with 70% and 30% probability (side and color counterbalanced).

Finally, we assessed learning. Participants were told, “Another kid played this game, and she clicked on the [blue/green] machine four times.” Participants selected how many candies they thought were received on those turns. This measure was completed twice for each machine, resulting in a “value estimate” for each machine ranging from zero of eight trials (0%) to eight of eight trials (100%).

Results

Children’s Engagement We first assessed whether parents reported any difficulty with the task. Only five out of 55 parents reported any technical difficulties, all minor. One parent reported influencing their child’s responses “once or twice,” while the remaining reported no interference. Fifty-three parents reported their children were somewhat or very engaged, and only one parent reported that their child took any breaks. In sum, the majority of children were engaged, provided their own responses, and encountered no technical problems.

Qualitative Predictions Next, we tested several qualitative predictions, which together allow us to assess whether participants understood the task and utilized the four-point curiosity scale as instructed. We tested these predictions separately for children and adults.

First, if participants attended to the observed outcomes, we would expect final value estimates to be higher for the 70% machine than the 30% machine. We fit a mixed-effects regression model predicting participants’ value estimates as a

function of machine, with random intercepts for participant. We tested the effect of machine using a likelihood ratio test. Children provided higher value estimates for the 70% machine ($M = 0.69$, $SD = 0.22$) than for the 30% machine ($M = 0.52$, $SD = 0.29$), $\chi^2(1) = 571.89$, $p < .001$, as did adults (70% machine: $M = 0.66$, $SD = 0.20$; 30% machine: $M = 0.43$, $SD = 0.24$), $\chi^2(1) = 1116.70$, $p < .001$.

Second, if participants understood the hidden agent and the curiosity rating scale, curiosity ratings should be low after an outcome that could not have been caused by AJ: candy in the thief condition, and no candy in the donor condition. We fit a mixed-effects regression model predicting curiosity ratings, with fixed effects for outcome, agent condition, and their interaction and with random intercepts for participant. We compared this model to a reduced model excluding the interaction, revealing a significant interaction in both children, $\chi^2(1) = 88.25$, $p < .001$, and adults, $\chi^2(1) = 472.65$, $p < .001$. Curiosity followed the predicted pattern in both groups: participants were less curious after receiving candy (vs. no candy) in the thief condition, and vice versa in the donor condition.

Triggers of Curiosity Next, we tested whether curiosity tracked surprise, uncertainty, and expected learning. To test whether each feature was associated with curiosity in isolation, we fit several mixed-effects regression models predicting curiosity ratings for adults or children. Each model included one feature as a fixed effect. Following Liquin et al. (2020), curiosity ratings were z-scored within participants, to account for variation in use of the rating scale. We also included by-participant random intercepts in each model. The results of this analysis are presented in Table 1. For both children and adults, the strongest predictor of curiosity was query EIG/uncertainty, closely followed by value EIG.

Surprisingly, surprise-IT and surprise-full were negatively associated with adults' and children's curiosity, and surprise-MAP was negatively associated with children's (but not adults') curiosity. The former associations could be explained by the negative correlations between query EIG/uncertainty and surprise-IT/surprise-full (see Fig. 1). On trials where query EIG/uncertainty is high, unlikely outcomes can be partially explained away by the possibility of intervention and thus are less surprising. On trials where query EIG/uncertainty is low, unlikely outcomes are less likely to be attributed to the intervention and are thus more surprising. Therefore, if curiosity is positively related to query EIG/uncertainty, it will be negatively related to surprise-IT and surprise-full. This is not the case for surprise-mean and surprise-MAP, which do not account for the intervention.

Though this analysis reveals the strongest single predictors of curiosity, it does not disentangle these predictors. For example, value EIG might be related to curiosity merely by virtue of its strong association with query EIG/uncertainty ($r = 0.72$, see Fig. 1). To disentangle the influence of each feature, we fit several multiple regression models. Because the candidate surprise features were highly correlated (or highly correlated when controlling for other features), we fit a separate model for each surprise feature. Each model also included fixed effects for value uncertainty, query EIG/uncertainty, value EIG, age group, and the interactions between age group and each feature, as well as by-participant random intercepts. We label these models by the included surprise feature: IT model, full model, mean model, or MAP model.

All four models were nearly equivalent in AIC: IT model 5674, full model 5676, mean model 5678, and MAP model 5676. Because the models achieved comparable fit to the data, we analyzed the coefficients of all four models.




Figure 2: Regression coefficients predicting curiosity in children and adults, with 95% confidence intervals. Coefficients and confidence intervals are estimated from multiple regression models fit within each age group. Asterisks indicate statistical significance of interactions between age group and features in the cross-age models; * $p < .05$; ** $p < .01$; *** $p < .001$.

First, to test whether any features were differentially predictive of curiosity in children versus adults, we tested whether each interaction term was significant using likelihood ratio tests. The interaction terms between age group and surprise, value uncertainty, and query EIG/uncertainty were significant in all models (see Fig. 2). To further probe these interactions, we fit separate regression models to children's responses and adults' responses, allowing us to investigate the association between each feature and curiosity within each age group (see Fig. 2). Children's curiosity was negatively related to surprise across models, such that curiosity was higher after outcomes that were less surprising (when uncertainty and expected learning were held fixed). Adults' curiosity, in contrast, was unrelated to surprise in all models but the IT model, where it was positively related to surprise. Across models, children's curiosity was positively related to value uncertainty, while there was no evidence that adults' curiosity tracked value uncertainty. Both children's and adults' curiosity tracked value EIG, but children's curiosity did not track query EIG/uncertainty when controlling for some definitions of surprise.

In summary, when the candidate triggers of curiosity were disentangled, we found that children's curiosity was higher when expected learning was higher (holding surprise and uncertainty fixed), when uncertainty was higher (holding surprise and expected learning fixed), and when surprise was lower (holding uncertainty and expected learning fixed). Adults' curiosity was higher when expected learning was higher (holding surprise and uncertainty fixed), but there was little evidence for an additional effect of surprise or uncertainty. All four definitions of surprise produced similar multiple regression results, suggesting similar associations with curiosity despite subtly different definitions.

Discussion

By precisely quantifying and disentangling candidate triggers of curiosity, we found that children's curiosity, like adults' curiosity, tracks expected learning—above and beyond heuristic cues (namely surprise and uncertainty). However, we also found evidence for developmental change. Whereas adults' curiosity was triggered almost exclusively by expected learning, children's curiosity was also related to surprise and uncertainty, tracking expected learning less closely.

What accounts for this developmental change? One possibility is that with increasing age, the elicitors of curiosity genuinely shift from heuristic cues (that may be easier to compute) towards "optimal" cues (that best track expected learning). However, it is also possible that this pattern of results reflects other developmental changes. First, the Bayesian model from which we estimated surprise, uncertainty, and expected learning could provide a better account of adults' learning than children's learning. For example, children might not fully integrate beliefs about the task's causal structure (e.g., the possibility of AJ intervening) into their learning (Cohen et al., 2020). Our model-estimated features capture surprise, uncertainty, and expected learning according to the optimal learner—and to the extent that children (or adults)

are not well-approximated by the optimal learner, these features may not be appropriate predictors of children's or adults' curiosity. Future work would benefit from testing alternative models that make different assumptions about children's and adults' learning.

In addition, children and adults might have different goals—for example, there may be a shift from exploration to exploitation (Gopnik, 2020), or from guiding learning to guiding action. Such a shift might be accompanied by a change in curiosity itself: from a non-instrumental drive, as posited by classic theories (Loewenstein, 1994), to a motivational state more integrated with considerations of utility (Dubey & Griffiths, 2020). In the context of our task, a learner could aim to learn the probability of reward from each machine (as assumed by our model), or instead the best machine to *choose*. The "optimal" trigger of curiosity will depend on what the learning goal is, so future research might explore additional learning goals.

In addition to considering other sources of developmental change, future work would benefit from investigating developmental change within childhood. When do the elicitors of curiosity change from more "child-like" to more "adult-like," and why? And is there an earlier period in development where children *only* track surprise and uncertainty?

Finally, the finding that children are more curious after outcomes that are less surprising (controlling for uncertainty and expected learning) is unexpected in light of prior research. It is possible that children's curiosity about unsurprising outcomes is specific to our task. Indeed, the surprising outcomes in our task were statistically unlikely, but they did not violate principles of core knowledge or intuitive theories (Bonawitz et al., 2012; Stahl & Feigenson, 2015), and thus only captured a low to moderate range of surprise. However, it is also possible that children's responses to surprise in prior research in fact reflect a preference for uncertainty or expected learning, which were confounded with surprise. If this is the case, a surprise that does not generate uncertainty or expected learning would not elicit curiosity (see Perez & Feigenson, 2020)—and might actually depress curiosity. Unsurprising observations could elicit more curiosity because they provide a unique opportunity to confirm one's beliefs, aligning with children's use of a "positive test strategy" during exploration—choosing questions or actions that are expected to provide confirmatory evidence for one's working hypothesis (Nussenbaum et al., 2020; Ruggeri et al., 2016). Future research could shed light on these issues.

In sum, we found both continuity and change in the triggers of curiosity across development. Although expected learning appears to be an important driver of curiosity in both children and adults, children show more sensitivity to other features and track expected learning less closely. Our findings raise important questions about children's information search: why do children attend, explore, and ask questions when they do? We suggest that surprise, uncertainty, and expected learning may all be at play—but these features must be disentangled to fully understand self-directed learning, including how it changes across the lifespan.

Acknowledgements

This work was supported by an NSF Graduate Research Fellowship to EL [DGE 1656466]. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Abir, Y., Marvin, C., Geen, C. van, Leshkowitz, M., Hassin, R., & Shohamy, D. (2020). Rational curiosity and information-seeking in the COVID-19 pandemic. *PsyArXiv*.
- 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.
- Coenen, A., Nelson, J. D., & Gureckis, T. M. (2019). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic Bulletin & Review*, 26(5), 1548–1587.
- Cohen, A. O., Nussenbaum, K., Dorfman, H. M., Gershman, S. J., & Hartley, C. A. (2020). The rational use of causal inference to guide reinforcement learning strengthens with age. *Npj Science of Learning*, 5, 16.
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, 120(3), 341–349.
- Dorfman, H. M., Bhui, R., Hughes, B. L., & Gershman, S. J. (2019). Causal inference about good and bad outcomes. *Psychological Science*, 30(4), 516–525.
- Dubey, R., & Griffiths, T. L. (2020). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, 127(3), 455–476.
- Fandakova, Y., & Gruber, M. J. (2021). States of curiosity and interest enhance memory differently in adolescents and in children. *Developmental Science*, 24(1), e13005.
- Gopnik, A. (2020). Childhood as a solution to explore-exploit tensions. *Philosophical Transactions of the Royal Society B*, 375, 20190502.
- Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49(10), 1295–1306.
- Jones, A., Swaboda, N., & Ruggeri, A. (2020). Developmental changes in question-asking. In L. P. Butler, S. Ronfard, & K. H. Corriveau (Eds.), *The Questioning Child: Insights from Psychology and Education*. Cambridge University Press.
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T., & Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, 20(8), 963–973.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS One*, 7(5), e36399.
- Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature Human Behaviour*, 3, 587–595.
- Liquin, E. G., Callaway, F., & Lombrozo, T. (2020). Quantifying curiosity: A formal approach to dissociating causes of curiosity. *Proceedings of the 42nd Annual Conference of the Cognitive Science Society* (pp. 309–315). Cognitive Science Society.
- Liquin, E. G., & Lombrozo, T. (2020a). A functional approach to explanation-seeking curiosity. *Cognitive Psychology*, 119, 101276.
- Liquin, E. G., & Lombrozo, T. (2020b). Explanation-seeking curiosity in childhood. *Current Opinion in Behavioral Sciences*, 35, 14–20.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1), 75–98.
- Lucca, K., & Wilbourn, M. P. (2018). Communicating to learn: Infants' pointing gestures result in optimal learning. *Child Development*, 89(3), 941–960.
- Nussenbaum, K., Cohen, A. O., Davis, Z. J., Halpern, D. J., Gureckis, T. M., & Hartley, C. A. (2020). Causal information-seeking strategies change across childhood and adolescence. *Cognitive Science*, 44(9), e12888.
- Perez, J., & Feigenson, L. (2020). Violations of expectation trigger infants to search for explanations. *PsyArXiv*.
- Poli, F., Serino, G., Mars, R. B., & Hunnius, S. (2020). Infants tailor their attention to maximize learning. *Science Advances*, 6(39), eabb5053.
- Rothe, A., Lake, B. M., & Gureckis, T. M. (2018). Do people ask good questions? *Computational Brain & Behavior*, 1(1), 69–89.
- Ruggeri, A., Lombrozo, T., Griffiths, T. L., & Xu, F. (2016). Sources of developmental change in the efficiency of information search. *Developmental Psychology*, 52(12), 2159–2173.
- Ruggeri, A., Sim, Z. L., & Xu, F. (2017). “Why is Toma late to school again?” Preschoolers identify the most informative questions. *Developmental Psychology*, 53(9), 1620–1632.
- 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.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379–423.
- Sim, Z. L., & Xu, F. (2017). Infants preferentially approach and explore the unexpected. *British Journal of Developmental Psychology*, 35(4), 596–608.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230), 91–94.
- Stojić, H., Orquin, J. L., Dayan, P., Dolan, R. J., & Speekenbrink, M. (2020). Uncertainty in learning, choice, and visual fixation. *Proceedings of the National Academy of Sciences*, 117(6), 3291–3300.
- Vogl, E., Pekrun, R., Murayama, K., & Loderer, K. (2020). Surprised–curious–confused: Epistemic emotions and knowledge exploration. *Emotion*, 20(4), 625–641.