

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Try smarter, not harder: Exploration and strategy diversity are related to infant persistence

Permalink

<https://escholarship.org/uc/item/2sc9m2cf>

Journal

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

ISSN

1069-7977

Authors

Radovanovic, Mia
Solby, Hannah
Soldovieri, Antonia
et al.

Publication Date

2021

Peer reviewed

Try smarter, not harder: Exploration and strategy diversity are related to infant persistence

Mia Radovanovic (m.radovanovic@mail.utoronto.ca)
Hannah Solby (hannah.solby@mail.utoronto.ca)
Antonia Soldovieri (antonia.soldovieri@mail.utoronto.ca)
Jessica A. Sommerville (jessica.sommerville@utoronto.ca)
Department of Psychology, University of Toronto

Abstract

Much research into persistence focuses on methods to increase trying without distinguishing whether persistence is rational. However, expectations of effort efficiency suggest that reducing effort in the face of repeated failure is logical. We performed archival behavioral coding to propose exploration as a rational means to extend persistence as new information is gained and the possibility of success is maintained. Infants were presented with an impossible task and exploratory behavior was classified. Infants decreased effort with increased experience failing, but persisted for longer when using several exploratory strategies and exploring for proportionally longer. These results confirm that infants are sensitive to the utility of their actions, and that exploration offers a means to persist even in the face of failure.

Keywords: infancy; persistence; exploration; utility maximization; effort conservation

Introduction

Colloquial speech is riddled with meritocratic messages encouraging constant persistence in the face of difficulty. However, energy and time are limited, and exerting sustained effort without payoffs is costly. If one is on a camping trip and seeking to find food, it is sensible to fish for an hour even without a catch. However, after an entire unsuccessful day, one ought to explore alternatives. This example illustrates a simple principle: persisting in the same action does not necessarily support broader goal-attainment.

Much of research into persistence begins with the assumption that increased persistence is always desirable or beneficial. This body of work suggests that persistence is a stable trait of individuals, rather than a contextual response, and links this disposition to outcomes both in childhood and throughout the lifespan. For example, grit is defined as an individual's perseverance and passion for long-term goals, and trait grit has been linked to greater academic performance in military cadets, as well as greater retention in high school, workplace sales, and even marriage (Duckworth et al., 2019; Eskreis-Winkler et al., 2014). Similar conclusions have been drawn about infants. For example, findings from Banerjee and Tamis-LeMonda (2007) suggest that infant persistence at 6 months uniquely predicts cognitive ability at 14 months.

While success would be impossible without sometimes enduring difficulty, there are drawbacks to wanton persistence, especially when task success is impossible or

when completion needs to be prioritized. To this end, Lucas et al. (2015) presented participants with a timed task in which unsolvable questions were interspersed with solvable questions. Participants higher in grit were less successful in completing the task compared to participants lower in grit, even losing monetary rewards, as they persisted on unsolvable questions. Thus, success often requires an understanding of when persistence ought to be deployed and when effort is better conserved for other tasks.

This is because effort is costly, requiring metabolic resources and creating inherent opportunity costs. Children display this intuition when observing others as young as 6 months old, expecting others to achieve goals through the most efficient paths possible (Liu & Spelke, 2017; Scott & Baillargeon, 2013). These expectations about effort efficiency are so strong that they have been formalized into a unified framework under the naive utility calculus (Jara-Ettinger et al., 2016), illustrating that children use effort expenditure as a cue to infer a wide range of information about others, including desires and competence (Jara-Ettinger et al., 2015), as well as social cooperation (Jara-Ettinger, Tenenbaum, & Schulz, 2015).

These inferences pertain to third-party evaluations, but children also seem to calibrate the effort they expend to its expected gain. Lucca, Horton, & Sommerville (2020) presented 18-month-olds with an impossible task and manipulated the level of success an adult demonstrator experienced prior to infant test trials. When the adult struggled and succeeded, infants assumed the task was difficult but their effort worthwhile, increasing trying time across trials and maintaining low levels of frustration. However, when the adult struggled and failed, infants assumed the task was unsolvable, expending consistently low effort and exhibiting higher frustration. These results suggest infants consider how worthwhile their efforts will be when deciding whether to persist, balancing the expected probability of success and the cost of effort.

Children's considerations of expected gain may not only pertain to task success, but also information gain. In Rett and Walker (2020), 4- and 5-year-old children attempted an impossible task, and some children were told that they would be given a solution at the end. When the solution was promised, children's efforts would not produce any necessary information that they later would not gain. However, without the promise of a solution, their effort was the only source of

knowledge. Accordingly, children persisted for longer and explored more without a promised solution, extracting the information that is available even in the absence of success.

These findings highlight the special role that information gain may play in the process of rational persistence. When there is uncertainty in problem-solving, it may pique curiosity and indicate opportunities to find solutions. If this uncertainty is paired with opportunities to engage, persistence is the natural result of independent exploration. There is a long history documenting the nuanced skills that infants have for independent exploration (Piaget & Inhelder, 1966; Sim & Xu, 2017). In particular, infants engage longer with tasks when expectations are violated. For example, Stahl and Feigenson (2015) presented infants with videos of objects violating various physical properties (e.g., solidity, continuity, support). Following observation of violations, infants explored objects for longer and learned more about associated properties. Importantly, infants explored the object strategically, allocating their efforts specifically to disambiguate the mechanism causing the event (e.g., dropping for violations of physical support). These findings suggest that surprising information does not only increase interest in objects, but also that energy is expended through exploration to systematically reduce uncertainty.

The present study seeks to provide an empirical evaluation of exploration in the process of persistence by extending an earlier experiment with infants through archival behavioral coding. The original experiment provided infants with experimenter demonstrations of a task, varying in success, then provided infants the opportunity to attempt the task independently. Here, we contrast imitation of the experimenter (a single action used to persist) with problem exploration to understand how allocation of time across strategies early in task trials related to overall persistence. We hypothesized that children as a group would prioritize exploration, as they could quickly obtain evidence that the demonstrated solution was ineffective. We also coded the number of strategies explored and the extent to which they varied from the experimenter demonstration. While the expected probability of success through imitation should converge towards zero with each failed attempt, new strategies should hold a stable, non-zero probability of success as they have not yet been tested. Therefore, as children diversify their problem-solving attempts, rather than doubling down with greater intensity on a single strategy, they should generate additional avenues to remain engaged rather than choosing to give up. Thus, we further hypothesized that exploring more, and attempting a greater number of strategies, would predict greater overall engagement.

Methods

Participants

Archival behavioral coding was performed using previously collected and published data by Lucca et al. (2020). To allow for sufficient, codable data, participants whose test trials did

not last for at least 30 seconds ($n = 25$) were excluded from our sample. This resulted in a final sample of 71 full-term, typically developing 18-month-olds recruited from a university database (mean age = 18.45 months, range = 17.67–19.23 months, 26 females). Their parents identified as either White ($n = 51$), mixed race ($n = 15$), Asian ($n = 3$), Hispanic ($n = 1$), or declined to report ($n = 1$).

Procedure

In the original study, infants were presented with a novel means-end task and took turns with an experimenter to attempt to solve the task. The experimental session was divided into four components: 1) warm-up, 2) experimenter demonstration, 3) infant test trials, and 4) a recovery trial. During the experiment, infants sat in their parent's lap. Between phases, parents faced away from the table towards a wall with pictures to distract infants. Parents wore occluding glasses and were instructed not to interfere.

To ensure infants' later performance was not due to shyness or disinterest, infants were presented with three novel toys during a warm-up and encouraged to play. Data was excluded from infants that did not play with at least two toys ($n = 5$). Next, infants were introduced to the means-end task by first watching an experimenter demonstration. Difficulty was manipulated using between-subjects conditions, such that the experimenter sometimes succeeded and sometimes failed. Unlike Lucca et al. (2020), we were only interested in individual differences in approach, thus, data were collapsed across conditions. During demonstrations, the experimenter attempted to retrieve an out-of-reach toy by pulling straight back, in the z -axis, on a rope connected to the toy (see Figure 1).

To measure persistence, infants completed three test trials, each after an experimenter demonstration. Success on the task was always impossible for infants. During test trials, the experimenter did not look at the infant unless the rope fell outside of the infant's reach, in which case the experimenter returned it within the infant's reach. Test trials were terminated when the infant did not interact with the rope for 15 seconds or after 120 seconds had passed.

Finally, to understand whether infants generalized their failure from the test trials, a recovery trial followed the three test trials. There was no experimenter demonstration before the recovery trial and success was possible. If the infant did not interact with the rope after 15 seconds, the experimenter prompted the infant and provided a hint. The experimenter continued prompting every 15 seconds until the infant retrieved the toy, or for a maximum of nine hints. The number of hints provided was used in subsequent analyses.

Behavioral Coding

Behavioral coding was performed to 1) identify exploratory behavior, 2) classify exploratory diversity and extent, and 3) catalogue overall engagement in the test trials. Coding began when the infant first grabbed the rope or when the parent's chair was still, whichever occurred first. Because infant's attempts were generally stochastic, coding was performed on

the level of 5-second intervals to allow for consistent classification of the primary behavior.

To ensure maximum independence between our measures of exploration and overall engagement, exploration was only coded during the first six intervals (i.e., 30 seconds), while overall engagement reflected how long infants engaged out of 24 possible intervals (i.e., 120 seconds). This decision also allowed us to code an equal amount of exploration data for all participants, as trials were terminated after 15 seconds of inactivity and there was substantial drop-off after the first 30 seconds (i.e., 96% of trials lasted at least 30 seconds, but only 78% of trials lasted at least 45 seconds).

Two independent coders were trained to apply the coding scheme to pilot data. This decision ensured independence between coders, as neither coder was instructed on intervals from the experimental data set. Once trained, each coder was assigned approximately half of the data set. In addition to coding assignments, approximately 25% of the sample was double-coded (16 of 71 participants' videos) to allow for comparison between coders.

Identifying Exploratory Behavior For each 5-second interval, the primary behavior which occurred was classified into five categories. There were two categories of theoretical interest: 1) exploration, and 2) imitation. Our primary interest in imitation was as a proxy for persistence using a single action, rather than as a social phenomenon. Thus, imitation was narrowly defined as single pulls in the z-axis, as a single straight-back pull was the action (repeatedly) taken by the experimenter in the demonstration. Exploration was defined as several behaviors which innovated on this approach (see Figure 1). Behavior was also divided into three off-task categories: 3) help-seeking, 4) play, and 5) nothing. These classifications (i.e., play, imitation, exploratory strategies) encompassed the full range of pulling behaviors infants enacted on the rope. Interrater reliability was excellent across all behavior types (*all a's* > .92). The measures used for analyses were the proportion of intervals spent imitating in the beginning of test trials and the proportion of intervals spent exploring in the beginning of test trials.

Because our data is correlational, we wanted to ensure maximum independence between our measures. One concern associated with our coding approach is that infants who had a greater disposition towards persistence would have a higher proportion of exploration or imitation in the first six trials because they were simply more active. If this were the case, then we would naturally expect these infants to persist longer overall. For this reason, we calculated an activity measure from the proportion of the first six intervals spent engaged (i.e., exploring, imitating, help-seeking, or playing). This activity measure was used to control for confounding associated with activity levels in later analyses.

Exploratory Extent and Diversity Because imitation and exploration were enacted on the same item (i.e., the rope), coding was also performed to identify the extent to which exploration deviated from the demonstration. For each



Exploratory Behavior	Classification
Examining while holding rope/handle	Examination
Single, extreme pull in the y- or x-axis	Micro
Episodic tugs in the z-axis	Micro
Adjustment of hand position	Micro
Bracing against table	Micro
Episodic, extreme pulls in the y- or x-axis	Macro
Pulling with hands-over-hands	Macro
Non-social reach directly for toy	Macro

Figure 1: Demonstration was performed in the z-axis, while exploration occurred in the x-axis and y-axis.

exploratory interval, coders classified the primary behavior into three categories: examination (cognitive exploration), macro-strategies (large deviations), and micro-strategies (slight deviations). Further, all unique exploratory strategies were tallied to generate a measure of exploratory diversity (see Figure 1). Interrater reliability was excellent for extent and diversity measures (*all a's* > .91).

Overall Engagement Because test trials were terminated after 15 seconds of inactivity, test trial length varied (range = 6-24 intervals; *M* = 16.54 intervals). To measure overall task persistence for each trial, coders were instructed to count the total number of intervals fully completed by the participant. Coders agreed on 100% of instances.

Results

Our first goal was to examine how infants allocated their time between exploration and imitation. We hypothesized that infants would prioritize exploration over imitation, as exploration represented several behaviors which had the possibility of “fixing” the impossible means-end problem, while imitation represented just one action. To answer this question, a paired two-sample t-test was used comparing early imitation and early exploration. Indeed, the proportion of the first six intervals spent exploring (*M* = 0.25, *SE* = 0.02) was greater than the proportion spent imitating (*M* = 0.11, *SE* = 0.01; *t*(212) = 6.75, *p* < .001; *d* = 0.46). Thus, while perhaps obvious, our hypothesis was supported.

Our next goal was to understand how exploration and imitation related to overall persistence. We hypothesized that

Table 1: Mixed effects models predicting overall engagement.

Model	AIC	Predictor					
Imitate	1340	<i>Random effects:</i>					
		Participant	Variance	10.50			
		<i>Fixed effects:</i>		Coefficient	SE	t-value	p-value
				Estimate			
		Intercept	16.101	1.689	9.534		
		Trial	-2.234	0.447	-4.997	< 0.001	
		Prop Active	5.775	1.460	3.955	< 0.001	
		Prop Imitated	4.383	2.600	1.688	0.093	
Explore	1339	<i>Random effects:</i>					
		Participant	Variance	9.777			
		<i>Fixed effects:</i>		Coefficient	SE	t-value	p-value
				Estimate			
		Intercept	16.246	1.673	9.710		
		Trial	-2.275	0.440	-5.165	< 0.001	
		Prop Active	5.063	1.535	3.299	0.001	
		Prop Explored	3.801	1.770	2.148	0.033	
Help-Seeking	1360	<i>Random effects:</i>					
		Participant	Variance	9.726			
		<i>Fixed effects:</i>		Coefficient	SE	t-value	p-value
				Estimate			
		Intercept	16.227	1.661	9.768		
		Trial	-2.163	0.442	-4.894	< 0.001	
		Prop Active	6.947	1.405	4.945	< 0.001	
		Prop Help-Seeking	-5.575	2.038	-2.736	0.007	
Unique Behaviors	1360	<i>Random effects:</i>					
		Participant	Variance	10.11			
		<i>Fixed effects:</i>		Coefficient	SE	t-value	p-value
				Estimate			
		Intercept	21.018	1.167	18.007		
		Trial	-2.839	0.437	-6.505	< 0.001	
		Unique Behaviors	0.916	0.328	2.792	0.006	

the total intervals engaged would rationally decrease across trials, as infants experienced repeated failure. We further predicted that greater exploration, but not greater imitation, would predict more intervals engaged overall. We built two linear mixed-effects models predicting overall engagement entering participant as a random effect and the main effect of trial as a predictor (see Table 1). As our data was correlational, we wanted to ensure observed relations were not merely indicative of more active participants. To mitigate this concern, the main effect of activity (i.e., the proportion of the first six intervals spent engaged) was also entered as a predictor in both models. As our questions pertained to individual differences in approach, we did not enter condition as a predictor for the sake of parsimony. However, it is worth noting that no condition effects were significant when

behavioral variables were entered and that the observed pattern of results did not vary when condition was included in models.

In one model, the proportion of the first six intervals spent imitating was entered as a third main effect. As hypothesized, the total intervals completed declined across trials ($t(157) = -5.00, p < .001, \beta = -2.23, SE = 0.45$). Also as hypothesized, the proportion of the first six intervals spent imitating did not significantly predict the total intervals completed ($p = .09$). In the second model, the proportion of the first six intervals spent exploring was entered as a third main effect. As hypothesized, the total intervals completed once again declined across trials ($p < .001$). However, even when controlling for overall activity, the proportion of the first six intervals spent exploring predicted more intervals engaged

overall ($t(209) = 2.15, p = .03, \beta = 3.80, SE = 1.77$). Therefore, exploration uniquely related to persistence such that greater exploration early in test trials was associated with greater engagement overall, while persisting in a single imitative action was not.

These results regarding exploration may be due to the expectation that when trying new actions, success is still possible, whereas previously failed actions are unlikely to suddenly work. To this end, an exploratory correlation was performed comparing a participant's average proportion of exploration in the first six intervals to the number of hints needed during recovery trials. Indeed, greater exploration was significantly correlated with needing fewer hints during recovery trials ($r(68) = -0.33, p = .005$). On the other hand, a participant's average proportion of imitation in the first six intervals did not significantly correlate with decreased hinting during recovery ($p = .09$). Therefore, there is initial evidence that exploration prolongs persistence by maintaining an expectation of success even after failure.

Our third goal was to understand the nature of children's exploration and how it related to overall engagement. Specifically, we hypothesized that testing a greater number of unique, exploratory actions would lead to greater persistence overall. To this end, a final mixed effects model was built entering participant as a random effect and the main effects of trial and the number of unique behaviors as predictors (see Table 1). The number of unique exploratory behaviors predicted total intervals completed ($t(210) = 2.79, p = .005, \beta = 0.92, SE = 0.33$). This again suggests that the value of exploration may be in extending available avenues for persistence that do not require repeating failed action.

Finally, we wished to address an exploratory question regarding the extent that exploration deviated from imitation in prolonging overall persistence. To answer this question, we conducted correlations between strategy types and total intervals engaged, adjusting for multiple comparisons using the Holm method (see Figure 2). Interestingly, the two categories of physical strategies (i.e., micro-strategies and macro-strategies) did not relate to longer engagement (both p 's $> .10$). On the other hand, the proportion of the first six intervals spent engaged in examination did significantly relate to the total intervals completed ($r(211) = .23, p = .004$). Therefore, it seems that studying the problem may enable children to think of more ways to engage. To confirm this possibility, a final correlation was performed comparing the number of unique strategies employed and the proportion of the first six intervals spent examining. Indeed, greater time examining was related to testing more unique strategies ($r(211) = .37, p < .001$). Altogether, this suggests that exploration may be an active cognitive and physical process which enables greater overall persistence through information search and the possibility of success.

Discussion

To solve challenging problems, infants must persist through difficulty. However, persistence is labor-intensive, and effort should be conserved when there is evidence of diminishing

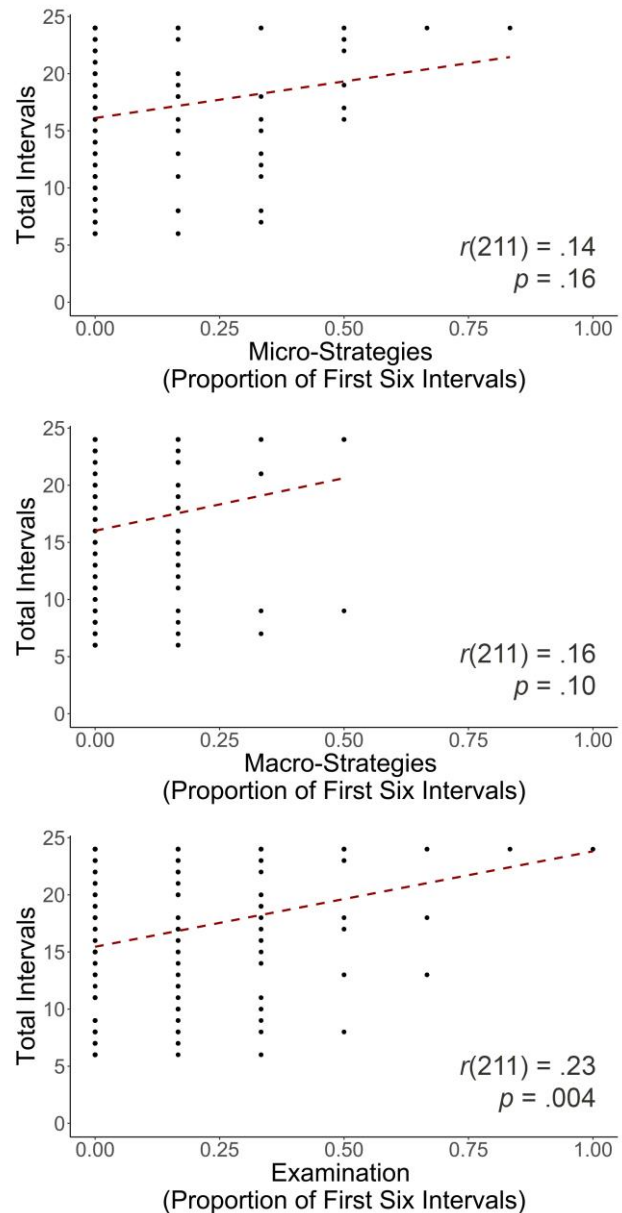


Figure 2: Correlations between strategies and total intervals completed, corrected for multiple comparisons.

returns. Diversification of problem-solving through exploration may offer a means of buffering motivational losses and extending effort. To understand whether infant persistence is sensitive to principles of effort conservation, and if exploration indeed provides a rational means to prolong engagement, we coded infant trying behaviors during an impossible task. In relation to our objectives, our results show that (a) infants prioritized exploratory behaviors over single-faceted, imitative behaviors, (b) that infants' persistence decreased with increased evidence of failure across trials, but that exploration extended persistence even when controlling for general activity, (c) similarly,

employing a greater diversity of exploratory strategies extended persistence, and (d) examination may have facilitated diversified problem-solving attempts and greater overall persistence.

For our first objective, we examined how infants distributed their time across imitation and exploration. Behavioral coding revealed that as a group, infants spent more time exploring than imitating. These results build on previous literature by demonstrating that infants' problem-solving attempts are not only sensitive to utility but calibrated in nuanced ways. At the group level, infants were discerning in how they allocated their time, and persisted adaptively by producing actions which could potentially resolve uncertainty rather than prioritizing single-dimensional, previously failed actions.

These results may seem intuitive given our single-dimensional definition of imitation. However, it is precisely because there are varied and numerous ways to explore that exploration would be a valuable problem-solving tool. To this end, our second objective was to understand how engagement varied on an individual level as a function of trial number and time spent engaged in exploration and imitation, respectively. Once again, infants conserved their efforts by reducing overall persistence with greater evidence of failure across trials, consistent with prior literature (Lucca et al., 2020; Sommerville et al., 2018).

While there is only one way to imitate, increased time imitating would still reflect greater determination and activity levels, and thus, engagement. However, imitation did not relate to overall persistence. On the other hand, greater time exploring in the beginning of trials did predict greater persistence even when controlling for activity levels. Therefore, infants divested from imitation at the group level, and did not seem to be aided by increased imitation on the individual level. Conversely, participants who explored more on average required less support during recovery trials than participants who explored less. Together, these results provide initial evidence that exploration helps maintain a belief that success is possible and may in turn increase overall persistence.

Further, our findings regarding unique exploratory strategies also support a conception of exploration as enabling persistence through undiscovered possibilities. Participants that attempted a greater number of exploratory strategies also demonstrated greater levels of persistence overall. Thus, it seems that exploration may enhance persistence precisely because participants are moving away from unsuccessful strategies to untested strategies, which may produce useful information or generate a successful result. This in turn could prolong the period of time in which engagement is rational and perceived as worthwhile.

Finally, we sought to explore the role of different strategy types. Interestingly, only the proportion of time spent examining in the beginning of trials related to overall persistence, and not the time spent engaging in various physical strategies. We followed this finding with an additional exploratory analysis, finding that greater time

examining related to a greater number of unique exploratory strategies being tested. Thus, these findings suggest a picture of exploration as an integrated cognitive and physical process. They also lend credibility to the idea that exploration did not increase engagement merely because it represented greater time manipulating items, but rather, because it offered a rational avenue to remain engaged with the problem even in the face of repeated failure.

Of course, this work is correlational, so it is unclear exactly what the causal relationship between exploration and persistence is. For example, it may be that more persistent infants also explore more as a function of dispositional differences. These dispositional differences may also serve to unify convergent measures (e.g., expectations during recovery). Further experimental work should provide the required verification to make the assertion that exploration itself drives greater persistence, as well as to link the structural pathways between examination, motivational variables, exploration, and time spent persisting. However, we did take multiple measures to mitigate the role of confounding factors in our data, both by restricting overlap between our measures of persistence and exploration, and by controlling for general activity levels in our analyses. Even with these controls, exploration related to persistence through multiple, converging relationships.

These findings open several questions. For example, one limitation of our study is that the task was truly impossible to solve so that particular strategies (i.e., prioritizing exploration, choosing to give up) were adaptive within our problem-solving context. However, problems are often solvable but still quite challenging. It may be the case that different strategies (i.e., high-fidelity imitation) are more productive in these contexts. Likewise, much research into persistence focuses on the ways in which parental language can enhance children's mindsets (Cimpian et al., 2007; Gunderson et al., 2013; Lucca & Sommerville, 2018). Perhaps, language encouraging greater exploration and divergent thinking would help children navigate feeling "stuck" on difficult problems. Likewise, exploratory priming may extend overall engagement by encouraging children to find new possibilities to test. Most interestingly, this work leaves an open question about whether the primary benefit of exploration is in its relationship to potential information gain or potential success. There may be individual variance in focus on these two outcomes, which may also drive persistence and learning disparities.

While much prior research has been devoted to motivating increased persistence in the face of failure, encouraging repetition of previously failed actions is not likely to lead to success in real situations. If we truly want to motivate children to endure difficulty, our research should create frameworks with real promise of gain. Indeed, we find evidence that infants conserve their energy when failing, but that exploration may support persistence by offering possibility and supporting this gain, both in the form of success and novel information. In this way, exploration may offer an organic means to extend persistence.

Acknowledgements

We would like to thank Kelsey Lucca and Rachel Horton for their work that has made this project possible. We would also like to thank the families and infants that participated in this work. This work was supported by grants from the Canadian Foundation for Innovation and the Ontario Research Foundation.

References

- Banerjee, P. N., & Tamis-LeMonda, C. S. (2007). Infants' persistence and mothers' teaching as predictors of toddlers' cognitive development. *Infant Behavior and Development, 30*(3), 479-491.
- Cimpian, A., Arce, H.-M. C., Markman, E. M., & Dweck, C. S. (2007). Subtle linguistic cues affect children's motivation. *Psychological Science, 18*(4), 314-316.
- Duckworth, A. L., Quirk, A., Gallop, R., Hoyle, R. H., Kelly, D. R., & Matthews, M. D. (2019). Cognitive and noncognitive predictors of success. *Proceedings of the National Academy of Sciences, 116*(47), 23499-23504.
- Eskreis-Winkler, L., Shulman, E. P., Beal, S. A., & Duckworth, A. L. (2014). The grit effect: Predicting retention in the military, the workplace, school and marriage. *Frontiers in Psychology, 5*.
- Gunderson, E. A., Gripshover, S. J., Romero, C., Dweck, C. S., Goldin-Meadow, S., & Levine, S. C. (2013). Parent praise to 1- to 3-year-olds predicts children's motivational frameworks 5 years later. *Child Development, 84*(5), 1526-1541.
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences, 20*(8), 589-604.
- Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Children's understanding of the costs and rewards underlying rational action. *Cognition, 140*, 14-23.
- Jara-Ettinger, J., Tenenbaum, J. B., & Schulz, L. E. (2015). Not so innocent: Toddlers' inferences about costs and culpability. *Psychological Science, 26*(5), 633-640.
- Liu, S., & Spelke, E. S. (2017). Six-month-old infants expect agents to minimize the cost of their actions. *Cognition, 160*, 35-42.
- Lucas, G. M., Gratch, J., Cheng, L., & Marsella, S. (2015). When the going gets tough: Grit predicts costly perseverance. *Journal of Research in Personality, 59*, 15-22.
- Lucca, K., Horton, R., & Sommerville, J. A. (2019). Keep trying!: Parental language predicts infants' persistence. *Cognition, 193*, 104025.
- Lucca, K., Horton, R., & Sommerville, J. A. (2020). Infants rationally decide when and how to deploy effort. *Nature Human Behaviour, 4*(4), 372-379.
- Piaget, J., & Inhelder, B. (1966). *The psychology of the child*. Presses Universitaires de France.
- Rett, A., & Walker, C. M. (2020). Knowing when to quit: Children consider access to solutions when deciding whether to persist. *Proceedings of the 42st Annual Conference of the Cognitive Science Society, 2877-2883*.
- Scott, R. M., & Baillargeon, R. (2013). Do infants really expect agents to act efficiently? A critical test of the rationality principle. *Psychological Science, 24*(4), 466-474.
- Sim, Z. L., & Xu, F. (2017). Learning higher-order generalizations through free play: Evidence from 2- and 3-year-old children. *Developmental Psychology, 53*(4), 642-651.
- Sommerville, J. A., Enright, E. A., Horton, R. O., Lucca, K., Sitch, M. J., Kirchner-Adelhart, S. (2018). Infants' prosocial behavior is governed by cost-benefit analyses. *Cognition, 177*, 12-20.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science, 348*(6230), 91-94.