

Infants Track Environmental Volatility to Optimize Their Learning

Francesco Poli (francesco.poli@mrc-cbu.cam.ac.uk)
University of Cambridge, Herchel Smith Bldg, Robinson Way
Cambridge, CB2 0SZ UK

Tommaso Ghilardi (t.ghilardi@bbk.ac.uk)
Birkbeck University of London, Malet Street
London WC1E 7HX UK

Jana Bersee (jana.bersee@student.uva.nl)
University of Amsterdam, Amsterdam Science Park 904
Amsterdam, 1098 XH the Netherlands

Rogier B. Mars (rogier.mars@ndcn.ox.ac.uk)
University of Oxford, John Radcliffe Hospital
Oxford, OX3 9DA UK

Sabine Hunnius (sabine.hunnius@donders.ru.nl)
Donders Institute, Thomas Van Aquinostraat 4
Nijmegen, 6525GD the Netherlands

Abstract

Infants' bodies, brains, and environments are ever-changing. Although this continuous transformation is a fundamental feature of development, how infants actively adapt and learn amidst such volatility is still unknown. To address this, we devised a novel learning task in which the location of a reward was systematically altered, transitioning from stable to volatile periods. Through computational modelling, we inferred from the infants' gaze and pupil data the learning processes that enabled them to navigate these changing environments. We found that infants' tonic pupil size reflected trial-by-trial changes in the level of environmental volatility. Moreover, phasic changes in pupil size when observing the reward indicated that infants relied on the information about volatility to optimize their learning. This resulted in the successful performance of the task, as indicated by the pattern of anticipatory looks to the correct reward locations. Together, these results identify the active role that infants play in adapting to change.

Keywords: Infancy; Learning; Computational modelling; Pupilometry; Volatility.

Introduction

As infants develop, their perceptual, motor, and cognitive skills undergo a rapid transformation, continuously altering their experience of the world (Hunnius, 2022; Westermann et al., 2007). For example, after lying supine during the first months of their lives, infants gain the ability to sit independently, which revolutionises how they interact with the environment and learn from it (Karasik et al., 2011). Cognitive development is marked by alternating periods of stability and volatility, presenting infants with the ubiquitous challenge of adapting to change.

Beyond these internal developmental dynamics, infants are also confronted with volatility stemming from external

events (Walasek et al., 2022). The effects of volatile environments on cognitive development have been documented extensively (Frankenhuis et al., 2016; Li et al., 2023). In volatile environments, children display higher vigilance (Silvers et al., 2017), are more exploitative (Xu et al., 2023), act more impulsively (Peviani et al., 2019), and favour instant over delayed rewards (Kidd et al., 2013). Recent work shows that volatile environments do not only impact immediate behaviour, but also affect brain development, changing the brain's structure and connectivity (Carozza et al., 2023) with a cascade of potential consequences for later cognitive development and psychosocial wellbeing (Davis et al., 2017; Molet et al., 2016).

Yet, the active role that infants might play in adapting to volatile environments remains unexplored. Both internal and external sources of change may require infants to estimate the degree of volatility that they are facing, and adjust their learning accordingly. Stable environments allow for the maintenance of learned behaviours, whereas increased volatility necessitates the discontinuation of past behaviours in favour of new adaptations. Research in adults has shown that humans can estimate environmental volatility to improve their learning efficiency (Behrens et al., 2007; Vincent et al., 2019). Moreover, variations in this ability have been linked to multiple mental conditions, to the extent that (mis)estimation of volatility has been proposed as a transdiagnostic process that broadly affects mental health (Sandhu et al., 2023). Hence, understanding how infants develop the ability to estimate and adapt to volatility may offer critical insights into the early mechanisms that shape cognitive resilience and vulnerability (Moscarello & Hartley, 2017).

In this study, we employed a novel experimental paradigm and computational modelling approach (Figure 1) to identify the early origins of volatility estimation, shedding light on the active role played by infants in responding to environmental change. First, we demonstrate that infants are capable of actively adjusting their behaviour in response to volatility, rather than being at the mercy of (and reactive to) external events. Second, we investigate the cognitive mechanisms that allow infants to achieve this learning. Infants might learn by constantly disregarding previous evidence, only valuing current events. This would allow them to respond relatively well to the ongoing challenges, but at the cost of storing little or no information from the past – a phenomenon known as catastrophic forgetting (Kirkpatrick et al., 2017; Zosh & Feigenson, 2015). Alternatively, infants might learn by adjusting their learning rate depending on the level of volatility, thus displaying an optimal form of learning.

Methods

Participants

Infants (N = 61, age = 7.8 months, SD = 0.3, F = 29) were recruited via a database with volunteer families. Infants who were born prematurely or had visual impairments were excluded from the recruitment. During the task, trials were presented until the infant lost attention or became fussy. For 6 infants, data was absent due to fussiness or lack of calibration. In addition, infants with 80% or more missing trials were excluded (N=17) resulting in a final sample of 38 infants. Caretakers of participating infants received either 10 Euros or a children’s book as compensation. The study was approved by the faculty’s board of ethics.

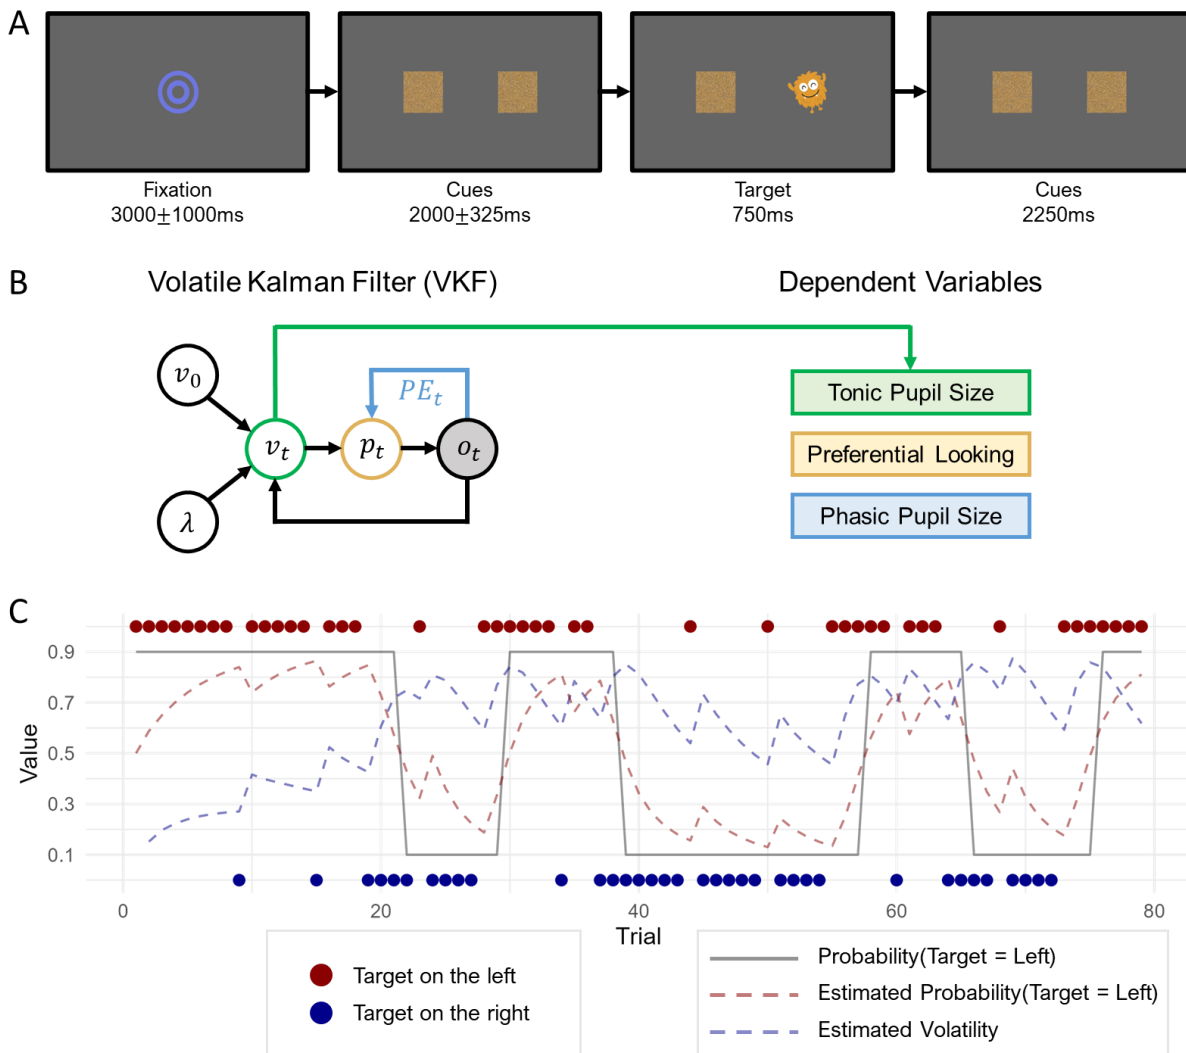


Figure 1: An example trial from the reversal learning task (A). A schematic representation of the volatile Kalman filter (VKF) and the dependent measures collected from infants’ gaze and pupil data, where v_0 indicates the initial volatility, λ is the volatility learning rate, v_t is the volatility of each trial, p_t is the belief about the target location, and o_t is the observed outcome of the target location (B). The target locations order (red and blue dots) with the VKF’s probability estimates about the target location (red line) and the VKF’s volatility estimates (blue line) (C).

Procedure

Caretakers were asked not to distract or redirect the infant's attention during the task. Infants were tested in a quiet room without daylight. The infant was positioned in front of the screen, either directly on the caretaker's lap or in a baby seat at a distance between 60 and 65 centimetres from the eye-tracker. After a 5-point calibration, the reversal learning task started. During the task, gaze and pupil data were collected using a Tobii X300 eye-tracker with Python via the `tobii_research` module.

Materials

Each trial consisted of a fixation bullseye, two cue boxes, and a target, which were presented on a screen at 1920x1080 pixel resolution. All images were 250x250 pixels and were presented on a grey background (#656565). The fixation bullseye was presented in the centre of the screen and the cues and target were presented 300 pixels to the left or the right of the centre. The target stimuli consisted of 40 different fantasy figures. The cues were created by scrambling the target stimuli and reshaping them into a square. This way, the luminance of the target and the cue were kept constant.

Experimental Paradigm

As depicted in Figure 1A, all trials started with the fixation bullseye (3000 ms \pm 1000 ms). Then, the cues were presented on the left and right side at the same distance from the center of the screen (2000 ms \pm 325 ms). At the start of the cue presentation, a sound was played to signal to the infants the start of a new trial. During the cue presentation, the cues rotated with an angle of 20 degrees and an accompanying sound was played to attract the infants' attention. Afterwards, a target stimulus appeared in place of one of the two cues for 750 ms. The target was also accompanied by a sound. After the target presentation, the same cues were presented statically for 2250 ms.

Four different sequences were generated and they were presented to participants in a pseudo-randomized order. An example of a sequence of trials can be seen in Figure 1C. The target appeared in a high-likelihood location approximately 90% of the times, and in a low-likelihood location the remaining 10% of the times. The high-likelihood location (e.g., left) remained stable for the initial 18 trials. Then, it changed for 9 trials (e.g., right), and again for 9 more trials (e.g., left). Afterwards, the high-likelihood location returned to stable for 18 trials (e.g., right). This stable-changing pattern allowed us to introduce variation in volatility levels, with stable periods being less volatile, and changing periods being more volatile. The stable-changing pattern was repeated until infants stopped looking at the screen for more than one minute or became fussy. The task was programmed in Python 3.6 using PsychoPy software (Peirce et al., 2019).

Measures

From the pupillometry data, we extracted a measure of tonic pupil size during the fixation period (i.e., before the trial

started) and a measure of phasic pupil size during the target presentation. Phasic pupil size was baseline-corrected using the 500ms preceding the target stimulus presentation. In research with adults, tonic pupil size has been shown to reflect subjective uncertainty (Muller et al., 2019), while phasic pupil size tracks the amount of information contained in a stimulus (Zénon, 2019) and whether such information is used to improve future predictions (O'Reilly et al., 2013).

From the gaze data, we extracted the proportion of anticipatory looking to each cue location before the target appeared. This measure is widely used as an index of infants' expectations (Téglás & Bonatti, 2016) and it was used here to assess whether infants correctly predicted the most likely target location. Specifically, the proportion of anticipatory looking indexed the looking time spent over the left cue, divided by the overall looking time for both left and right cue. We expect this value to be high when infants are anticipating that the target is on the left, and low when they expect it to be on the right.

Computational modelling

A binomial volatile Kalman filter (Piray & Daw, 2020) was employed to track trial-by-trial changes in volatility (Figure 1B). The model learns the most likely target location via a trial-by-trial updating rule:

$$p_t = p_{t-1} + \alpha_t(o_t - s(p_{t-1}))$$

where the probability of where the target will appear (i.e., left or right side) is updated depending on the previous belief p_{t-1} , adjusted by the prediction error (i.e., the difference between the actual outcome o_t and the previous belief). The previous belief is mapped to the unit range (i.e., [0,1]) via a sigmoid function s . The prediction error is weighted by the learning rate α_t , which is computed as follows:

$$\alpha_t = \sqrt{w_t + v_t}$$

where w_t is the variance of the probability p_t , which can be seen as first-order uncertainty (i.e., uncertainty in the outcome), while v_t is the volatility of the environment, which can be seen as second-order uncertainty (i.e., uncertainty in whether the environment will change). Volatility is updated on every trial with its own update rule:

$$v_t = v_{t-1} + \lambda(\Delta v_t)$$

where λ is the learning rate that determines how much volatility prediction errors Δv_t change previous expectations about volatility v_{t-1} and was introduced in the model as a free parameter. The value of v_t at the start of the task (i.e., when $t = 0$) was introduced as an additional free parameter v_0 . The free parameters λ and v_0 were fitted based on the tonic pupil size during the fixation period, such that the relation between volatility and tonic pupil size was maximized with a grid-search algorithm:

$$Pupil_t = \beta_0 + \beta_1 v_t + \beta_2 t + f(\text{time}, \text{infant}) + \varepsilon$$

where tonic pupil size on each trial was predicted by volatility estimates v_t , trial number t (which allowed us to control for changes in pupil dilation over time), a smoothing function that captured constant fluctuations across time within the trial (in milliseconds) and across infants, as well as noise ε .

Results

Infants track environmental volatility

The model estimates of environmental volatility significantly correlated with infants' tonic pupil size ($t = 139.17$, $\beta = 0.23$, $SE = 0.001$, $p < .001$), indicating that infants were successfully tracking environmental volatility (Fig. 2A). Specifically, the model parameters that best predicted infants' phasic pupil size were λ (learning rate) = .30 and v_0 (initial volatility) = .01. These values indicate that infants started the task expecting a stable environment (i.e., low initial volatility), and successfully adapted their beliefs about environmental volatility as the environment changed, as indicated by levels of λ different from zero.

The model predictions about the most likely target location significantly correlated with infants' anticipatory looking ($t = 3.78$, $\beta = 1.68$, $SE = 0.44$, $p < .001$) (Fig. 2B). This indicates that, by tracking environmental volatility, infants flexibly adjusted their predictions about where the target was most likely to appear. Crucially, infants successfully predicted the target locations not only in stable but also in volatile environments. This demonstrates that infants were not simply more uncertain or confused when volatility was high, but were instrumentally using the information about environmental volatility to optimize their learning.

Infants optimize their learning

To shed light on the mechanisms that underlie the optimisation of learning in changing environments, we analysed the relation between infants' phasic pupil size and the trial-by-trial prediction errors estimated by the VKF model. In volatile environments, high prediction errors likely signal a change in the observed regularities, and should thus have a large impact on the observer's existing expectations. Conversely, in stable environments, high prediction errors are likely to be isolated instances that should be disregarded.

Consistent with these predictions, the interaction between volatility and the magnitude of the prediction errors significantly modulated phasic pupil size ($t = 14.71$, $\beta = 0.10$, $SE = 0.006$, $p < .001$). As depicted in Fig. 3, when the environment was more volatile, greater prediction errors led to a greater phasic pupil response ($t = 10.01$, $\beta = 0.3$, $SE = 0.003$, $p < .001$). This indicates that their importance was up-weighted. Conversely, when the environment was more stable, greater prediction errors led to a reduction in phasic pupil size ($t = -16.20$, $\beta = -0.05$, $SE = 0.003$, $p < .001$). This indicates that their importance was down-weighted. Hence, infants optimised their learning by flexibly weighting the impact of the prediction errors, depending on the volatility of the current environment.




Figure 2: Infants' tonic pupil size correlated with the VKF model's estimates of volatility (A). Infants' proportion of anticipatory looking correlated with the VKF model's predictions about the target location (B). Predictive means (y-axes) were obtained generating trial-by-trial estimates of pupil size and proportion of anticipatory looking from the beta coefficients of the fitted regression models.

Discussion

While previous work has addressed the impact of volatile environments on cognitive development, the active role that infants play in responding to environmental change is unexplored due to a lack of precise behavioral or physiological measures that probe into infants' learning abilities, as well as a lack of computational methods that can adapt to the higher levels of noise that are typical of developmental populations.

Here, we used a novel experimental paradigm that systematically manipulated volatility levels, and we quantified these trial-by-trial changes in volatility with computational modelling. By relating infants' tonic pupil size to the model's estimates of volatility, we showed that infants track changes in volatility, with their tonic pupil size reducing when the environment is more stable and increasing when the environment is more volatile.

As indexed by phasic pupil size in response to the target presentation, tracking volatility allowed infants to flexibly weight the relevance of incoming prediction errors, such that prediction errors were up-weighted when the environment was volatile and environmental change was likely, but down-weighted when the environment was stable and change was unlikely. In turn, this adaptive learning led to correct predictions of the target location both in stable and volatile periods, as indicated by the infants' anticipatory looking to the cue locations that were more likely to display the target.

This pattern of results indicates that infants are not only sensitive to changes in environmental volatility, but they

track and use this information about volatility to optimize their learning. As such, these findings promote an outlook on cognitive development where the infants and their adaptive skills play an active role in shaping the effects of the environment on their developing minds.

In adults, abnormalities in the ability to adapt to change have been linked to multiple mental conditions, to the extent that (mis)estimation of volatility has been proposed as a transdiagnostic process that broadly affects mental health (Sandhu et al., 2023). Revealing how infants estimate and adapt to volatility offers critical insights into the early mechanisms that shape vulnerability and resilience (Moscarello & Hartley, 2017). First, the stability of the environment (and lack of thereof) might affect the learning strategies that infants acquire: Infants who grew up in volatile environments might be able to respond promptly when the world is unpredictable, but struggle in decreasing their learning rate when stability is (re)introduced. Conversely, infants who only experienced stable environments might lack adaptiveness when new changes arise. Given the active role played by infants in adapting to change, cognitive development might not only depend on the early exposure to different kinds of environments, but also on the infants' own ability to adapt to them. From a simple change in routines to more fundamental changes in caregivers or social dynamics, it is the infant's ability to adapt to change that will ultimately shape their experience. Rigid infants might struggle even in more stable environments, while adaptive infants might be able to better tackle volatile situations. Hence, it is likely to be the combination of early experiences and the infants'

Phasic Pupil Size

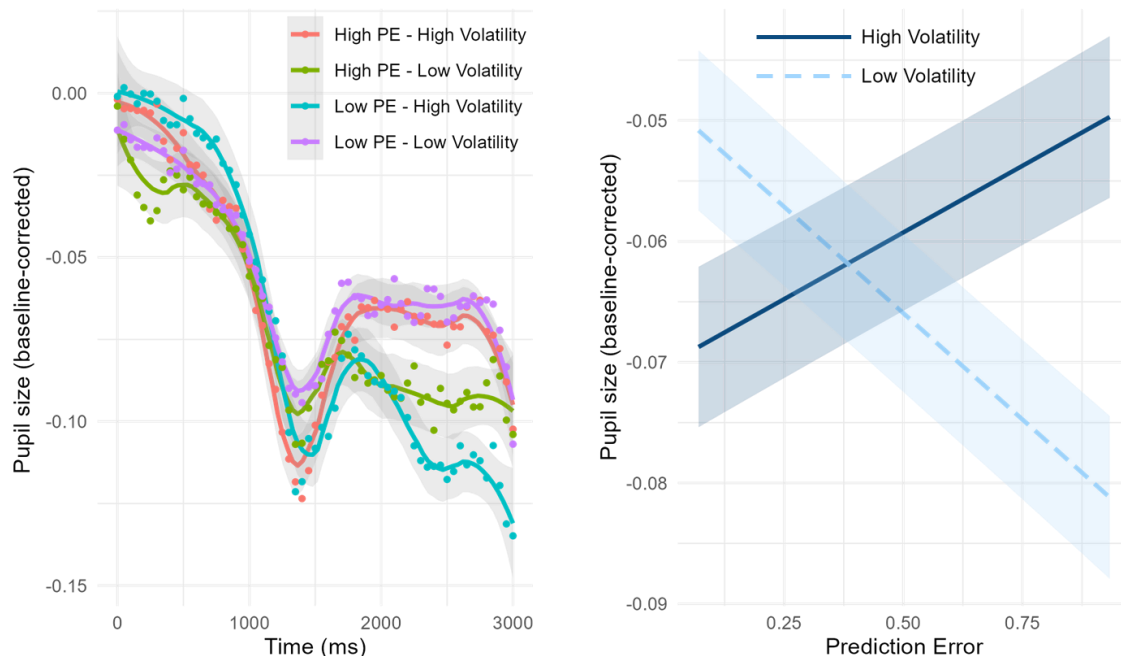


Figure 3: Raw data for phasic pupil size (baseline corrected) at the moment of the target presentation, divided in four groups depending on volatility (high or low) and the magnitude of the prediction error (big or small) (A). Predictive means of phasic pupil size as a function of environmental volatility and the magnitude of the prediction error (B). Phasic pupil size was significantly modulated by both volatility and prediction error, with greater levels of phasic pupil size for small prediction errors in more stable environments and for high prediction errors in more volatile environments.

adaptive learning abilities that together shape later psychosocial wellbeing. Research on individual differences in infants' volatility estimation – possibly across a variety of environments – is needed to ultimately understand how early variability in volatility (mis)estimation influences later cognitive development and psychological wellbeing.

References

- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214–1221. <https://doi.org/10.1038/nn1954>
- Carozza, S., Akarca, D., & Astle, D. (2023). The adaptive stochasticity hypothesis: Modeling equifinality, multifinality, and adaptation to adversity. *Proceedings of the National Academy of Sciences of the United States of America*, 120(42), e2307508120. <https://doi.org/10.1073/pnas.2307508120>
- Davis, E. P., Stout, S. A., Molet, J., Vegetabile, B., Glynn, L. M., Sandman, C. A., Heins, K., Stern, H., & Baram, T. Z. (2017). Exposure to unpredictable maternal sensory signals influences cognitive development across species. *Proceedings of the National Academy of Sciences of the United States of America*, 114(39), 10390–10395. <https://doi.org/10.1073/pnas.1703444114>
- Frankenhuis, W. E., Panchanathan, K., & Nettle, D. (2016). Cognition in harsh and unpredictable environments. *Current Opinion in Psychology*, 7, 76–80. <https://doi.org/10.1016/j.copsyc.2015.08.011>
- Hunnus, S. (2022). Early cognitive development: Five lessons from infant learning. In *Oxford Research Encyclopedia of Psychology*.
- Karasik, L. B., Tamis-LeMonda, C. S., & Adolph, K. E. (2011). Transition from crawling to walking and infants' actions with objects and people. *Child Development*, 82(4), 1199–1209.
- Kidd, C., Palmeri, H., & Aslin, R. N. (2013). Rational snacking: Young children's decision-making on the marshmallow task is moderated by beliefs about environmental reliability. *Cognition*, 126(1), 109–114. <https://doi.org/10.1016/j.cognition.2012.08.004>
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526. <https://doi.org/10.1073/pnas.1611835114>
- Li, Z., Sturge-Apple, M. L., Platts, C. R., & Davies, P. T. (2023). Testing different sources of environmental unpredictability on adolescent functioning: Ancestral cue versus statistical learning and the role of temperament. *Journal of Child Psychology and Psychiatry, and Allied Disciplines*, 64(3), 437–448. <https://doi.org/10.1111/jcpp.13714>
- Molet, J., Heins, K., Zhuo, X., Mei, Y. T., Regev, L., Baram, T. Z., & Stern, H. (2016). Fragmentation and high entropy of neonatal experience predict adolescent emotional outcome. *Translational Psychiatry*, 6(1), e702. <https://doi.org/10.1038/tp.2015.200>
- Moscarello, J. M., & Hartley, C. A. (2017). Agency and the Calibration of Motivated Behavior. *Trends in Cognitive Sciences*, 21(10), 725–735. <https://doi.org/10.1016/j.tics.2017.06.008>
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195–203. <https://doi.org/10.3758/s13428-018-01193-y>
- Peviani, K. M., Kahn, R. E., Maciejewski, D., Bickel, W. K., Deater-Deckard, K., King-Casas, B., & Kim-Spoon, J. (2019). Intergenerational transmission of delay discounting: The mediating role of household chaos. *Journal of Adolescence*, 72, 83–90. <https://doi.org/10.1016/j.adolescence.2019.03.002>
- Piray, P., & Daw, N. D. (2020). A simple model for learning in volatile environments. *PLOS Computational Biology*, 16(7), e1007963. <https://doi.org/10.1371/journal.pcbi.1007963>
- Sandhu, T. R., Xiao, B., & Lawson, R. P. (2023). Transdiagnostic computations of uncertainty: Towards a new lens on intolerance of uncertainty. *Neuroscience and Biobehavioral Reviews*, 148, 105123. <https://doi.org/10.1016/j.neubiorev.2023.105123>
- Silvers, J. A., Goff, B., Gabard-Durnam, L. J., Gee, D. G., Fareri, D. S., Caldera, C., & Tottenham, N. (2017). Vigilance, the Amygdala, and Anxiety in Youths With a History of Institutional Care. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 2(6), 493–501. <https://doi.org/10.1016/j.bpsc.2017.03.016>
- Vincent, P., Parr, T., Benrimoh, D., & Friston, K. J. (2019). With an eye on uncertainty: Modelling pupillary responses to environmental volatility. *PLOS Computational Biology*, 15(7), e1007126. <https://doi.org/10.1371/journal.pcbi.1007126>
- Walasek, N., Frankenhuis, W. E., & Panchanathan, K. (2022). Sensitive periods, but not critical periods, evolve in a fluctuating environment: A model of incremental development. *Proceedings. Biological Sciences*, 289(1969), 20212623. <https://doi.org/10.1098/rspb.2021.2623>
- Westermann, G., Mareschal, D., Johnson, M. H., Sirois, S., Spratling, M. W., & Thomas, M. S. C. (2007). Neuroconstructivism. *Developmental Science*, 10(1), 75–83. <https://doi.org/10.1111/j.1467-7687.2007.00567.x>
- Xu, Y., Harms, M. B., Green, C. S., Wilson, R. C., & Pollak, S. D. (2023). Childhood unpredictability and the development of exploration. *Proceedings of the National Academy of Sciences*, 120(49), e2303869120. <https://doi.org/10.1073/pnas.2303869120>
- Zosh, J. M., & Feigenson, L. (2015). Array heterogeneity prevents catastrophic forgetting in infants. *Cognition*, 136, 365–380. <https://doi.org/10.1016/j.cognition.2014.11.042>