

UC Irvine

UC Irvine Previously Published Works

Title

Interval Timing as a Computational Pathway From Early Life Adversity to Affective Disorders

Permalink

<https://escholarship.org/uc/item/9hs4r76b>

Authors

Harhen, Nora C
Bornstein, Aaron M

Publication Date

2023-10-12

DOI

10.1111/tops.12701

Copyright Information

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

Peer reviewed



Topics in Cognitive Science 00 (2023) 1–21

© 2023 Cognitive Science Society LLC.

ISSN: 1756-8765 online

DOI: 10.1111/tops.12701

This article is part of the topic “Best of Papers from the 20th International Conference on Cognitive Modeling,” Terrence C. Stewart (Topic Editor).

Interval Timing as a Computational Pathway From Early Life Adversity to Affective Disorders

Nora C. Harhen,^a Aaron M. Bornstein^{a,b}

^a*Department of Cognitive Sciences, University of California, Irvine*

^b*Center for the Neurobiology of Learning and Memory, University of California, Irvine*

Received 1 March 2023; received in revised form 26 September 2023; accepted 27 September 2023

Abstract

Adverse early life experiences can have remarkably enduring negative consequences on mental health, with numerous, varied psychiatric conditions sharing this developmental origin. Yet, the mechanisms linking adverse experiences to these conditions remain poorly understood. Here, we draw on a principled model of interval timing to propose that statistically optimal adaptation of temporal representations to an unpredictable early life environment can produce key characteristics of anhedonia, a transdiagnostic symptom associated with affective disorders like depression and anxiety. The core observation is that early temporal unpredictability produces broader, more imprecise temporal expectations. As a result, reward anticipation is diminished, and associative learning is slowed. When agents with such representations are later introduced to more stable environments, they demonstrate a negativity bias, responding more to the omission of reward than its receipt. Increased encoding of negative events has been proposed to contribute to disorders with anhedonia as a symptom. We then examined how unpredictability interacts with another form of adversity, low reward availability. We found that unpredictability's effect was most strongly felt in richer environments, potentially leading to categorically different phenotypic expressions. In sum, our formalization suggests a single mechanism can help to link early life adversity to a range of behaviors associated with anhedonia, and offers novel insights into the interactive impacts of multiple adversities.

Keywords: Early life adversity; Associative learning; Interval timing; Anhedonia

Correspondence should be sent to Nora C. Harhen, Department of Cognitive Sciences, University of California, Irvine, 2201 Social & Behavioral Sciences Gateway, Irvine, CA 92679, USA. E-mail: nharhen@uci.edu

1. Introduction

Across development, brain circuits adapt to reflect the environment's structure, preferentially encoding more frequent aspects of the world. The statistics of the early life environment tune sensory receptive fields, producing nonhomogeneous sensitivity to perceptual stimuli and determining discrimination abilities in adulthood (Efrati & Gutfreund, 2011; Tanaka, Ribot, Imamura, & Tani, 2006). Early consistency in these sensory inputs is crucial for the future functionality of involved circuits (Li, Fitzpatrick, & White, 2006). Similar developmental processes may take place in reward and memory systems, those underlying associative learning, implying that the consistency or predictability of associations in early life may shape the acquisition of associations later on (Birnie et al., 2020).

Caregivers are primary contributors to the associative structure infants encounter. Associations may take the form of a caregiver's response to an action the infant preforms. These responses may vary in their valence and predictability. Valence influences whether the infant will repeat the action preceding the response, while predictability constrains the infant's learning to associate the two. Prior work has largely focused on the effect of valence on later child mental health outcomes (Belsky & Fearon, 2002; Hane, Henderson, Reeb-Sutherland, & Fox, 2010; NICHD Early Care Research Network, 2006; Sroufe, 2005). However, recent work has highlighted how early life unpredictability, or ELU, may also contribute (Baram et al., 2012). Research done in animals has illustrated that offspring exposed to unpredictable caregiver signals show a reduction in motivation and the experience of pleasure, characteristics of the trans-diagnostic symptom, anhedonia (Bolton et al., 2018). Work in humans accords with these findings, showing the relationships between experiences of ELU, reduced reward anticipation, and symptom severity in anhedonia, depression, and anxiety (Dillon et al., 2009; Goff et al., 2013; Hanson et al., 2016; Mehta et al., 2010; Spadoni et al., 2022).

Here, we propose that the study of ELU can be understood in part via its influence on the development of temporal representations (TRs) that serve as basis sets for associative learning more generally (Howard et al., 2014; Jin, Fujii, & Graybiel, 2009). TRs capture the intuition that the strength of learned associations is dependent on the time between events (Balsam, Drew, & Gallistel, 2010). These tuning curves are similar to those found in sensory areas, but rather than being tuned to visual angle or auditory pitch, are sensitive to the temporal duration between related events.

We specifically examine how ELU can, via its influence on the adaptation of temporal representations, result in an anhedonic phenotype. We extend a principled computational model of interval timing (Ludvig, Sutton, & Kehoe, 2008) to simulate how enhanced volatility during an early period of heightened plasticity can, with minimal assumptions, affect later predictions of reward during maturity. With this model, we formally demonstrate that early unpredictability in timing, and adaptation of temporal representations to this timing, can lead to the development of several defining characteristics of anhedonia—including slowed associative learning, reduced motivation, and a bias toward learning from negative events—in the absence of differences in the overall amount of reward. Our results reproduce empirical findings that unpredictability in early life experience can heighten susceptibility to poor mental health

outcomes even after controlling for the childhood environment's overall resource availability (Glynn et al., 2019).

While we show that a singular type of adversity can alone produce an anhedonic phenotype, in the real world, individuals are often subject to multiple adversities. Modeling the nature of these interactions and their combined effect on learning will be critical for characterizing the developmental trajectory of psychopathology. As a first step, we model how temporal unpredictability interacts with the environment's availability of reward, or richness, to shape later learning and expectations of reward. Under the common cumulative risk approach to conceptualizing and measuring early life adversity (Felitti, 2002), these two adversities are assumed to have an additive effect on development: individuals facing both are predicted to have the most negative outcomes. Our model predicts that unpredictability always has a negative effect on associative learning; however, contrary to the cumulative risk prediction, this effect is most pronounced in richer environments. Both unpredictability and an abundance of rewards individually alter TRs to be more expansive or diffuse, producing the observed interaction. Our results highlight the potential value of computational psychiatric approaches to tackling the heterogeneity of early life adversity and making sense of its developmental consequences.

2. Isolating the contributions of one form of adversity, unpredictability

2.1. Methods

During the initial phase ("critical period"), agents' TRs were allowed to adapt to the environment's temporal statistics. Agents belonged to one of two groups, ELU or control. The two groups were differentiated by the distributions their reward timings were sampled from, with the ELU agents' distribution having the same mean as the control agents' but a higher variance. In the second phase ("post critical period"), both groups received reward at the same time step on each rewarded trial and, critically, agents' TRs were not allowed to adapt to the novel environment's statistics.

2.1.1. The temporal-difference learning model

Temporal-difference (TD) models aim to accurately estimate the value of world states, V , in terms of the future rewards they predict. Time is explicitly represented in these models with each time step identifying a world state.

$$V^* = E \left[\sum_{k=1}^{\infty} \gamma^{k-1} r_{t+k} \right] \quad (1)$$

where r_t is the reward received at the current time step, and γ is a parameter controlling how heavily future rewards are discounted. Future rewards are less influential on the estimation of V when γ is low. A TD agent learns V via an error-driven learning rule—the difference, δ_t , between the reward received ($r_t + \gamma V_t$) and the previously predicted reward (V_{t-1}) is used to




Fig. 1. Stimulus encoding by microstimuli. From left to right, the memory trace produced by a stimulus is approximated with a set of temporal basis functions, whose centers vary such that they evenly cover the trace's possible heights. The decaying nature of the memory trace produces microstimuli that become shorter and wider the further their center is to the stimulus onset. The microstimuli are weighted and averaged to estimate the future expected reward following the stimulus. The weights can be adjusted with experience to support accurate predictions of reward.

update the estimate of V at the next time step.

$$\delta_t = r_t + \gamma V_t - V_{t-1} \quad (2)$$

2.1.2. Microstimulus representation of time

All TD models explicitly represent time, but do so in various ways. Basic TD models use a complete-serial-compound (CSC) representation in which each time step is treated as independent from one another and agents are assumed to have perfect knowledge of when events occur. This representation prohibits temporal generalization, creating issues in environments where the time between cue and reward varies. The microstimulus representation addresses this problem by relaxing its temporal markers (Ludvig et al., 2008). CSC's discrete markers are replaced with continuous "microstimuli" which allow for temporal uncertainty to be represented (Fig. 1). A stimulus is assumed to leave behind a memory trace that decays with time. The trace is approximated by a set of Gaussian temporal basis functions uniformly distributed across the heights of the memory trace. This approximation produces a set of microstimuli increasing in their peak and width from the time of stimulus onset.

$$f(y, \mu, \sigma) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(y-\mu)^2}{2\sigma^2}} \quad (3)$$

A time step's value, V_t , is estimated as the weighted average of the microstimuli.

$$V_t = w_t^T x_t = \sum_{i=1}^n w_t(i) x_t(i) \quad (4)$$

V_t is compared to the reward received to compute an error term, δ_t that is used to adjust the weights on the microstimuli. Adjusting the weights updates the predicted value at the next time step.

$$w_{t+1} = w_t + \alpha \delta_t e_t \quad (5)$$




Fig. 2. Simulated agents learned to associate a cue with reward in two different environments. The cue was partially reinforced—75% of the time in the initial environment and 55% in the second. On rewarded trials, reward was delivered at a variable time step. Agents belonged to one of two groups, differing in the variability they experienced in the initial environment. The reward timings experienced by agents in the early life unpredictability (ELU) group were on average the same as those experienced by the control group. However, in the initial phase (“critical period”), they experienced more variably timed rewards trial to trial. In the second phase (“post critical period”), agents’ weights were frozen, and all agents received reward at the same time step.

α is the learning rate controlling the time window over which experiences are integrated. e_t is a vector containing each stimulus’s eligibility traces.

$$e_t = \gamma \lambda e_t + x_t \tag{6}$$

Following the stimulus, its eligibility trace decays at a rate determined by γ and λ . γ is a temporal discounting factor as it was for the TD model with a CSC representation, while λ controls the time window over which a stimulus can induce learning within a trial. For all simulations, we use the parameter settings from Ludvig et al. (2008)— $\alpha = 0.01$, $\gamma = 0.98$, $\lambda = 0.95$, $n = 50$, and $\sigma = 0.08$.

2.1.3. Simulating development

To model developmental changes in learning, we restrict the updating of microstimuli weights to the initial period which we treat as a “critical period” during which the TRs are tuned to support accurate estimation of V . This adaptation process is designed to mimic the observed tuning of sensory receptive fields during analogous sensitive periods of development (Simoncelli & Olshausen, 2001). During the second phase (“post critical period”), the weights are frozen, prohibiting representation adaptation, to simulate adulthood.

We simulated two groups of agents learning to associate a cue with reward across the two phases (Fig. 2). One group of agents, the ELU group experienced a volatile critical period environment in which the timing of reward was much more variable than the timing experienced by the control group. Critically, however, the average timing of reward and the

average amount of reward received (i.e., same probability of reward on each trial) was matched between groups.

On each of the 1000 simulated trials during the critical period, a cue was always presented at 10 time steps and there was a 75% probability of a reward following it. If a cue was followed by reward on a trial, the timing of reward was sampled from a normal distribution with $\mu = 30$ and truncated at 10 and 70 time steps. σ varied between agents. For agents in the ELU group, σ was sampled from a zero-truncated normal distribution with $\mu_{\text{hyper,elu}} = 10$ and $\sigma_{\text{hyper,elu}} = 3$. The control group experienced much less variability, with σ being sampled from a zero-truncated normal distribution with $\mu_{\text{hyper,control}} = 1$, $\sigma_{\text{hyper,control}} = 2$. We varied σ within groups to reflect the variation observed in real-life samples, particularly early life adversity facing ones, and to ensure our results were robust to such variation.

In the second phase, the microstimuli weights were frozen (“post critical period”), allowing us to directly examine the influence of highly variable early life experiences. The temporal statistics of this environment differed from the critical period’s environment in two ways: (1) The reward was delivered at the same time step each trial for both groups of agents. (2) This time step was later (50 time steps) than the previous environment’s average time of reward (30 time steps). By testing ELU agents’ learning in novel environments that are more stable than the environment they “developed” in, we formalize the Mismatch Hypothesis of Early Life Adversity and Depression (Schmidt, 2011). Under this hypothesis, depression and other mental illnesses are proposed to be the byproduct of a mismatch between the developmental environment to which neural systems are optimized for and the later adulthood environment. We were particularly interested in characterizing how an agent’s early adaptation to unpredictability would affect their response to uncertainty in adulthood. Within the simulated task, uncertainty should rise once the mean time of reward has passed and reward has failed to be delivered. This is because it becomes unclear whether the reward is delayed or is omitted altogether on the trial. To produce this circumstance, we moved back the time step of reward in the novel, post critical period environment to examine how the ELU and control groups differ in their response to reward and its omission following a period uncertainty. All agents completed two trials. On both trials, the cue arrived at 10 time steps. On one trial, reward followed the cue at 50 time steps. On the other, reward was omitted. We simulated agents only on two trials because the weights were no longer updated. Thus, the prediction error response on every trial of the same time (rewarded vs. omitted) would be identical.

2.1.4. Statistical analyses

Each simulated agent encountered a different sequence of reward timings during the initial critical period. Thus, a potential concern is that our results are largely driven by a subset of simulated agents. To assess the reliability of the relationship between prediction error magnitude and unpredictable experience, we performed a bootstrap analysis across agents within a group (Bornstein et al., 2023; Kim, Lewis-Peacock, Norman, & Turk-Browne, 2014). For each group, we sampled agents with replacement until we reached the total number of agents, 100. We then computed the test statistic for a two sample *t*-test with the selected groups. We repeated this procedure 1000 times to obtain a distribution of test statistics across shuffled permutations of the simulated groups. This resampling procedure provides a *p*-value that is




Fig. 3. (A, B) Positively weighted microstimuli. With experience, the ELU group grew to more heavily weigh delayed, imprecise microstimuli to account for the frequent delayed rewards. (C) **Temporal Imprecision.** We computed a summary statistic of temporal representation (TR) imprecision by taking a weighted average of the standard deviations of the positively weighted microstimuli at the end of the critical period. ELU agents' temporal representations were, on average, more than twice as imprecise as control agents.

the fraction of test statistic values with a different sign from the base effect size (the test statistic for the original two groups). We also computed the Cohen's d in order to evaluate the size of the difference between simulated populations. By convention, effect sizes greater than 0.80 are considered "large", and thus reliable (Cohen, 1992).

2.2. Results

2.2.1. Critical period

First, we validated that the critical period environment shaped TRs by comparing the groups' microstimuli weights at the end of the critical period. For each agent, we computed a temporal imprecision measure by taking a weighted average of the microstimuli's standard deviations, with the weights being the same as those used to generate the value signal. Consistent with our prediction that TRs would adapt to reflect the statistics of their environment, we found that the ELU group relied on more broadly tuned TRs relative to controls (Fig. 3; $t(198) = 8.43, p < .001, \text{Cohen's } d = 1.19$).

ELU has been shown to produce slower learning from reward in adulthood (Birn, Roeber, & Pollak, 2017; Dillon et al., 2009). We next examined the model's ability to capture this. As a proxy for learning, we used a particular pattern of prediction error responses. If a cue has become associated with reward, then there should be large positive prediction error in response to the cue, a smaller positive prediction error at the time of reward, and a large negative prediction error when reward is omitted. To compare prediction errors between groups, we computed, across time within each trial, the prediction error extremum for each agent. On rewarded trials, the maximum prediction error magnitude following the cue was taken and on omission trials, the minimum was taken. We then took the average of these




Fig. 4. An example ELU and control agent's prediction errors (δ) from individual trials within the critical period. A cue always occurred at 10 time steps, while the reward's timing varied from trial to trial. Temporal variability was determined by which group an agent belonged to—an ELU agent experienced a much wider distribution of reward times. Reward elicited a strong positive prediction error from both agents on the first trial. Even very early on, the control agent demonstrated a positive prediction error in response to the cue, a weak positive prediction error at the time of reward, and a strong negative prediction error when reward was omitted, matching the pattern of responses expected for well-learned, consistent contingencies using this temporal-difference learning rule. This pattern held throughout the 1000 trial critical period. In contrast, even very late into the critical period, the ELU agent's prediction errors continuously moved around in time and were larger in magnitude, a consequence of their more volatile environment.




Fig. 5. Critical period prediction error signals. Reward elicited larger positive prediction errors in ELU agents, while reward omission produced weaker negative prediction errors, a pattern of responses suggesting that ELU agents were slower in learning from reward.

values across trials of the same type for each participant. We found that, on rewarded trials, the ELU group's positive prediction errors were larger than the control group's (Figs. 4 and 5; $t(198) = 12.59$, $p < .001$, Cohen's $d = 1.78$) but, were less negative on omission trials ($t(198) = 6.23$, $p < .001$, Cohen's $d = 0.88$). Despite both groups experiencing the same amount of reward, the ELU group showed slower learning under reinforcement. Collectively, these results demonstrate how impaired associative learning, as observed in anhedonia, can emerge from experienced temporal volatility alone during a period of plasticity.

ELU has also been shown to impair motivation (Hanson, Williams, Bangasser, & Peña, 2021), potentially stemming from reduced expectations of reward. Thus, we next compared




Fig. 6. The value signal, V , averaged over all critical period trials. Individual agents' value signals are depicted by the thin lines. The thicker lines depict the group averages. Control agents' expectations of future reward quickly rose following the cue and steadily increased until the average time of reward, after which their expectations quickly dropped. ELU agents' expectations of reward similarly rose in response to the cue but subsequently decreased at a gradual rate rather than increasing. Notably, ELU agents had higher expectations of reward at later time steps compared to controls—a consequence of having experienced more delayed rewards which required relying on more diffuse, later peaking microstimuli. When aggregated across trials, ELU agents' expectations were more spread out. This is both because they relied on more diffuse microstimuli and because their value signals fluctuated from trial to trial in response to variably timed rewards.

the groups' expectations of future reward, as reflected by their value signals. When averaged across trials, control agents' value signals quickly increased in response to the cue (Fig. 6; mean at 10 time steps = 0.43, $sd = 0.022$), gradually rose until the average time of reward (mean at 26 time steps = 0.71, $sd = 0.075$) after which the signal rapidly dropped off (mean at 32 time steps = 0.059, $sd = 0.078$). ELU agents' value signals similarly rose in response to the cue but peaked much earlier ($t(198) = -27.75$, $p < .001$, Cohen's $d = -3.92$) and fell more gradually (mean at 32 time steps = 0.29, $sd = 0.045$, $t(198) = 26.34$, $p < .001$, Cohen's $d = 3.73$). Importantly, ELU agents' expectations of reward were diminished at the time steps right before when reward as most likely (mean at 26 time steps = 0.48, $sd = 0.048$, $t(198) = -25.87$, $p < .001$, Cohen's $d = -3.66$). These differences could have a particularly significant impact on decision-making which requires deciding not only *which* option to take but also *when* to take it. Diminished expectations of reward should produce slower decision times, a characteristic found in anhedonia (Dubal, Pierson, & Jouvent, 2000; Day et al., 2015; Gollan, Pane, McCloskey, & Coccaro, 2008; White, Myerson, & Hale, 1997). ELU agents also showed greater variability in their value signals from trial to trial as revealed by taking the standard deviation of the time steps at which value signals peaked (ELU mean = 10.49; Control mean = 1.50; $p < .001$, Cohen's $d = 3.39$). This aligns with prior empirical work that found more variable ventral striatal activity following early life stress (Hanson et al., 2016).

2.2.2. Post critical period

To simulate adulthood, in the second phase, we closed the "critical period" by preventing the updating of the microstimuli weights in the novel environment. Thus, their expectations

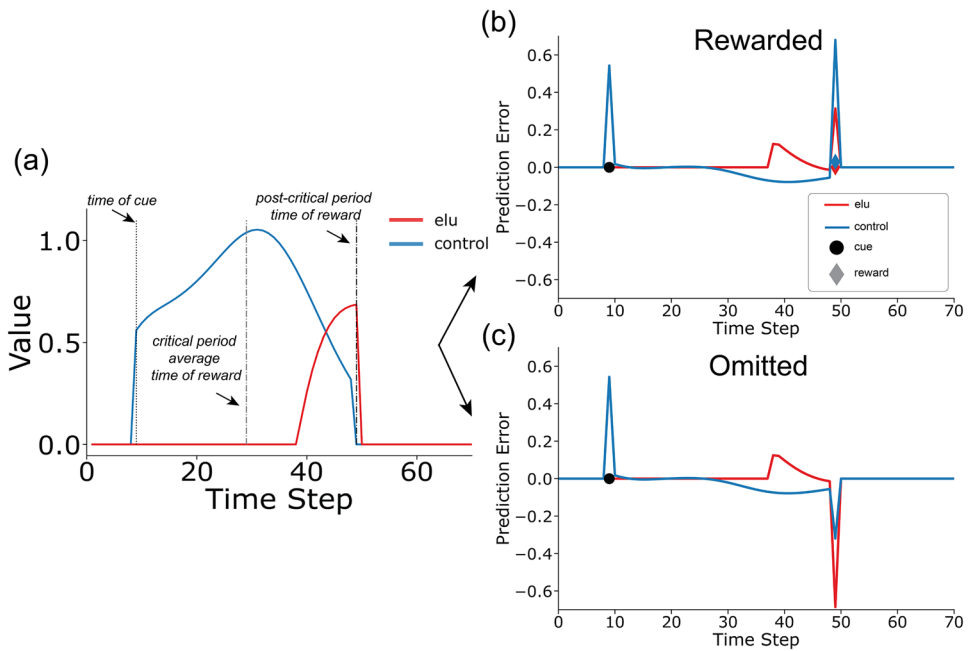


Fig. 7. (A) Representative agents' value signals. The value signal, taken from the end of the critical period, reflects the individual agent's expectation of future reward following the cue. These expectations are "frozen" and determine the agent's response to reward and its omission. (B) **Example prediction error signals for a single rewarded trial.** The ELU agent's expectation of future reward only begins to rise at 40 time steps, whereas the control agent's rises immediately at 10 time steps in response to the cue. Accordingly, the ELU agent demonstrates a weaker and delayed response to the cue. When reward is delivered at 50 time steps instead of its average previous time, 30 time steps, the control agent shows a more positive prediction error than the ELU agent. Again, this is a result of their expectations. The control agent does not expect the reward to arrive this late in the trial, and thus, is surprised when it does. The ELU agent, having experienced more delayed rewards, is less surprised. (C) **Example prediction error signals for a single omission trial.** The ELU agent's greater expectation of reward at later time steps also produces a larger negative prediction error when reward is omitted.

of reward are carried over and fixed once the developmental period ends. In this environment, reward was delivered at a later time step than the average time of reward during the critical period. This induces an interval of uncertainty during which it's unclear whether the reward is delayed or omitted. We examined how the expectations acquired in an unpredictable early life environment shape the prediction error response when this uncertainty is resolved. Because ELU agents experienced rewards at more variable time steps, they grew to have a higher expectation that reward could arrive at later time steps (Fig. 7A). This affects their response to the cue and reward. Control agents have a strong positive prediction error immediately after the cue is presented because they have learned well that the cue predicts reward (Fig. 7B). ELU agents instead have a weaker and delayed response to the cue because of their weaker association between the cue and reward. Control agents experience a slightly negative prediction error when reward is not delivered at the most expected time step (Fig. 7C).




Fig. 8. Sensitivity to increasing rewards. We varied the magnitude of rewards delivered during the second phase. As the magnitude of rewards increased, both groups showed larger positive prediction errors on rewarded trials. ELU agents were more sensitive to changes in reward magnitude—their prediction errors increased to a great extent in response to larger rewards. At the lowest reward magnitude, which was the magnitude experienced during the critical period, the control group experienced larger positive prediction errors than the ELU. This pattern reversed at larger magnitudes with ELU agents demonstrating hypersensitivity to rewards. Error bars are 95% bootstrapped confidence intervals.

But, when reward ultimately arrives at a later time step, they show a large positive prediction error, a consequence of their low expectations of reward this late in the trial. ELU agents had relatively higher expectations of reward at the time step when reward was delivered, thus they showed relatively blunted positive prediction errors ($t(198) = -2.25$, $p < .001$, Cohen's $d = -0.32$). The same expectations produced amplified negative prediction error when reward was omitted ($t(198) = -12.29$, $p < .001$, Cohen's $d = -1.74$). In other words, their higher expectations allowed them to experience greater disappointment.

We next examined how ELU affected agents' response to rewards of varying magnitudes. When given a reward of the same magnitude as those received during the critical period, control agents responded with larger positive prediction errors (Fig. 8; $\beta_{elc} = -0.51$, $p < .001$). As the reward magnitude increases, diverging from those previously experienced, both groups show increasingly large prediction errors ($\beta_{magnitude} = 0.55$, $p < .001$). The ELU agents do so at a faster rate than control agents, demonstrating larger prediction errors than controls in response to higher magnitude rewards ($\beta_{elc \times magnitude} = 0.43$, $p < .001$). When coupled with their blunted response to the cue, ELU agents appear to be hyposensitive to rewards in anticipation but hypersensitive to them in consumption. This pattern has been observed in a monetary incentive delay task designed to distinguish between reward anticipation and consumption (Boecker et al., 2014). More generally, it concurs with wide-spread findings that early life adversity impairs cue-reward learning (Birn et al., 2017; Dennison et al., 2019; Dillon et al., 2009; Stuart, Hinchcliffe, & Robinson, 2019) while increasing sensitivity to dopamine-releasing drugs (Cruz, Quadros, Planeta, & Miczek, 2008; Kosten, Miserendino, & Kehoe,




Fig. 9. Learning Asymmetry Indices. The ELU group showed a negativity bias, experiencing more extreme prediction errors on the omission trial than the rewarded trial. In contrast, the control group demonstrated a positivity bias, experiencing larger prediction errors on the rewarded trial. Error bars are 95% bootstrapped confidence intervals.

2000; Kosten, Sanchez, Zhang, & Kehoe, 2004; Kosten, Zhang, & Kehoe, 2006; Paine et al., 2021; Wakeford et al., 2019; Zhang, Sanchez, Kehoe, & Kosten, 2005).

Prediction error magnitude determines the extent to which an agent learns or updates their expectations. Because valence asymmetries in learning have been proposed to be clinically relevant (Pike & Robinson, 2022; Rouhani, Norman, Niv, & Bornstein, 2020; Rouhani & Niv, 2019), we next compared prediction error magnitude on the rewarded and omission trials to probe for such asymmetries. We computed an asymmetry index for each agent as follows:

$$index = \frac{|PE_+| - |PE_-|}{|PE_+| + |PE_-|} \quad (7)$$

ELU agents' asymmetry indices were overall negative (Fig. 9; $t(99) = -5.62$, $p < .001$, Cohen's $d = -0.79$), while the control agents' were positive ($t(99) = 8.49$, $p < .001$, Cohen's $d = 1.20$). Because prediction error magnitude enhances learning and memory, this suggests that negative events would have an outsized influence on ELU agents, making their value estimates overly pessimistic, while control agents' are overly optimistic (Sharot, 2011). Our model provides a mechanism through which both of these biases could emerge under minimal assumptions.

3. Interactions between multiple forms of adversity

3.1. Methods

3.1.1. Critical period

To examine the interaction between multiple forms of early life adversity—temporal unpredictability and low reward availability, we additionally manipulated the richness of the critical period environment and observed its effect on both groups. This allowed us to test the

assumptions of the cumulative risk conceptualization of early life adversity which assumes an additive effect of adversities on developmental outcomes. We simulated groups of the ELU and control agents in environments with 25% , 55% , 75% , and 95% probability of reward. As in previous simulations, the time of reward delivery was sampled from a normal distribution with $\mu = 30$ time steps and truncated at 10 and 70 time steps, and the distribution's σ differed between groups—ELU agents' σ were sampled from a zero-truncated normal distribution with $\mu_{hyper,elu} = 10$ and $\sigma_{hyper,elu} = 3$ and controls' were sampled from a zero-truncated normal distribution with $\mu_{hyper,control} = 1$, $\sigma_{hyper,control} = 2$.

3.1.2. Post critical period

In the novel environment during the second phase, the cue was presented at 10 time steps on each trial. They experienced one rewarded and one omission trial. On the rewarded trial, reward was delivered at 50 time steps. As before, we only include two trials because the weights are no longer updated, thus, the response on each trial of the same type would be identical.

3.2. Results

3.2.1. Critical period

As before, we assume that the smaller positive prediction errors are in response to reward and the larger negative prediction errors are in response to its omission then the more strongly an agent has learned to associate a cue with reward. Under this assumption, both temporal unpredictability and low reward availability were found to slow associative learning. On rewarded trials, positive prediction errors were larger for ELU agents and both groups' prediction errors became weaker with environment richness (Fig. 10A; $\beta_{elu} = 0.057$, $p < .001$, $\beta_{rich} = -0.90$, $p < .001$). On omission trials, negative prediction errors were stronger for control agents and with increasing environment richness (Fig. 10B; $\beta_{elu} = -0.022$, $p = .015$, $\beta_{rich} = -0.98$, $p < .001$). The two dimensions interacted, with the difference between groups increasing as environment richness increased ($\beta_{elu*rich} = 0.15$, $p < .001$). In particular, the effects of unpredictability on learning were only observed in richer environments, with no main effect of group but an interaction effect between group and richness ($\beta_{elu} = -0.015$, $p = .11$, $\beta_{elu*rich} = 0.093$, $p < .001$). Taken together, our results reveal that the effect of temporal unpredictability is most fully felt when reward is abundant, a consequence of both dimensions increasing the imprecision of TRs ($\beta_{elu} = 1.04e - 05$, $p < .0001$, $\beta_{rich} = 1.47e - 05$, $p < .001$, $\beta_{elu*rich} = -1.91e - 07$, $p = .95$). When rewards are both unpredictably timed and abundant, it increases the range of timings an agent's representation must accommodate.

The value signal reveals a similar impact of the environment's temporal unpredictability and overall richness on learning. The value signal correspondingly increased as richness increased (Fig. 10C; $\beta_{rich} = 0.32$, $p < .001$). Yet, only when the environment is sufficiently rich can unpredictability exert its blunting effect on the signal ($\beta_{elu} = 0.0041$, $p = .52$, $\beta_{elu*rich} = 0.035$, $p < .001$).

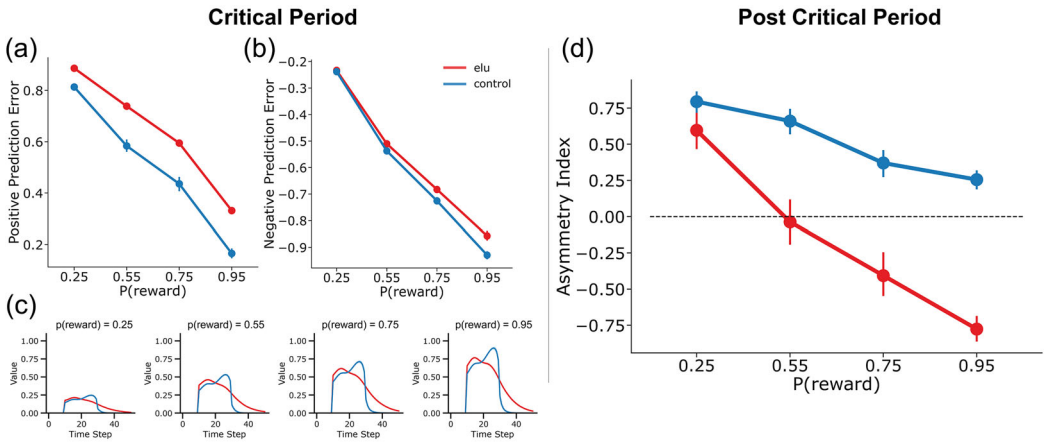


Fig. 10. Varying critical period environment richness to examine the impact of multiple adversities. **(A) Critical period prediction errors in response to reward.** Positive prediction error magnitude was modulated by the environment's richness (probability of reward) and its temporal unpredictability (ELU vs. Control), with richness attenuating magnitude and unpredictability amplifying it. **(B) Critical period prediction errors in response to reward omission.** Negative prediction error magnitude was amplified by richness and attenuated by unpredictability. This pattern of responding suggests that richness supports associative learning, while unpredictability impairs it. **(C) Value Signal.** Mirroring the reward statistics of their environment, agents' expectation of future reward increased accordingly with the overall richness of the environments. Notably, group differences were emphasized by richness. **(D) Post critical period asymmetry indices.** Control agents demonstrated a consistent positivity bias that diminished the richer the environment. ELU agents showed a positivity bias only in the poorest environment and a negativity bias in richer environments. Error bars are 95% bootstrapped confidence intervals.

3.2.2. Post critical period

In the post critical period phase, we found the same complex relationship between the environment's temporal unpredictability and richness, in which greater reward availability allows unpredictability to exert its influence. Across all environments, control agents maintained a bias toward learning from reward over its omission as indicated by positive asymmetry indices (Fig. 10D; 25% - $t(99) = 21.88$, $p < .001$, Cohen's $d = 3.09$; 55% - $t(99) = 15.79$, $p < .001$, Cohen's $d = 2.23$; 75% - $t(99) = 8.49$, $p < .001$, Cohen's $d = 1.20$; 95% - $t(99) = 8.62$, $p < .001$, Cohen's $d = 1.22$). The valence of ELU agents' biases, in contrast, was dependent on the richness of the critical period environment. ELU agents who experienced the sparsest rewards during the critical period exhibited a positivity bias, similar to control agents although weaker (Fig. 10D; 25% - $t(99) = 9.098$, $p < .001$ Cohen's $d = 1.28$). Those who experienced a less sparse environment showed no bias (55% - $t(99) = -0.46$, $p = 0.64$, Cohen's $d = -0.065$), and those who experienced an environment abundant with rewards exhibited a negativity bias (75% - $t(99) = -6.60$, $p < .001$, Cohen's $d = -0.79$; 95% - $t(99) = -17.72$, $p < .001$, Cohen's $d = -2.51$). This pattern of results is a byproduct of the reward expectations built up during the critical period. ELU agents whose representations are adapted for richer environments have a stronger prior expectation that reward will have a delayed arrival rather than being omitted altogether. Thus, when reward is omitted on a

trial, they experience a particularly large negative prediction error. Our simulations contradict the predictions that would be made under the cumulative risk approach which assumes an additive effect of adversities.

4. Discussion

Here, we propose a novel computational link between ELU and the emergence of anhedonia—the optimization of temporal representations (TRs) to the early life environment. By simply assuming that TRs are adapted to the statistics of the early life environment, several behaviors associated with anhedonia emerge—impaired learning from reinforcement, reduced anticipation of reward, and a greater response to the omission of events.

These findings are consistent with behavioral outcomes observed in the laboratory and clinical settings. One representative set of such findings is of an asymmetric attentional bias in anhedonia. If we treat the omission of reward as a negatively valenced event and the presence of reward as a positive event, this suggests a negative attentional bias in the ELU group and positive bias in the controls, reproducing empirical findings (Dillon & Pizzagalli, 2018; Frank, 2004). Larger negative prediction errors may not only affect attention in the moment but also have longer lasting consequences via memory. Surprising events, like prediction errors, are known to be more easily retrieved from memory (Rouhani et al., 2020; Sinclair & Barse, 2018). This provides a mechanism by which singular negative events can have an outsized influence on expectations and consequently, shape mood over the longer term (Eldar, Rutledge, Dolan, & Niv, 2016). Frequent large negative prediction errors could produce the persistent negative mood that characterizes anhedonia (Dillon et al., 2009). We found that the development of this negativity bias was critically dependent on the overall richness of the environment. To experience a pronounced negative prediction error when reward was omitted, agents needed to have a strong expectation that reward would come but a weak expectation of when that would be. Only in environments rich with variously timed rewards did such expectations emerge.

Our results contradict the assumptions and predictions of the cumulative risk conceptualization of early life adversity (Felitti, 2002). The cumulative risk approach has been crucial in establishing the robust association between negative events early in life and a wide array of negative outcomes later in development. However, aggregating over heterogeneous experiences may obscure the mechanisms linking such experiences to later psychopathology (McLaughlin, Sheridan, Humphreys, Belsky, & Ellis, 2021; Smith & Pollak, 2021). One proposed alternative are dimensional models which identify influential features of the early life environment on development and seek to characterize how these features exert their influence. Supporting the dimensional approach, recent work has found divergent associations between measures of threat and deprivation in the early life environment with later developmental outcomes, including amygdala reactivity to threat, aversive learning, cognitive control, and pubertal timing (Lambert, King, Monahan, & McLaughlin, 2017; Machlin, Miller, Snyder, McLaughlin, & Sheridan, 2019; Miller, Machlin, McLaughlin, & Sheridan, 2021; Rosen et al., 2020; Sheridan, Peverill, Finn, & McLaughlin, 2017; Sumner, Colich, Uddin,

Armstrong, & McLaughlin, 2019; Sun, Fang, Wan, Su, & Tao, 2020). However, adopters of these approaches have been criticized for an unprincipled choice of dimensions, particularly lacking neurobiological grounding (Smith & Pollak, 2021). Given the potential relevance of reward systems to psychopathology, it may be valuable to look at the statistical properties of the environment known to influence associative learning as potential candidate dimensions.

Thus far, in our interpretation of the results, we have treated the cue-paired outcome as reward. However, the model is agnostic to the valence of the outcome—allowing for different interpretations where the outcome is treated as neutral or aversive. Different valences will suggest different behavioral phenotypes. Treating the outcome as aversive, like a shock, the ELU group's prolonged expectation of a negative outcome's appearance could be interpreted as sustained hypervigilance (perhaps akin to a form of “paranoia”), a symptom of anxiety. Treating the outcome as neutral, impairments in associative learning become more general impairments in relational learning. This may explain memory deficits and alterations in hippocampal structure in ELU individuals (Granger et al., 2021; Molet et al., 2016) and anhedonia's associated memory deficits. Prior work has suggested that anhedonia is characterized not only by the inability to experience pleasure in the moment but also the inability to recall past and anticipate future pleasurable experiences (Dillon & Pizzagalli, 2018).

Here, we have only considered the mechanism under Pavlovian learning conditions. However, it also suggests differences in ELU individuals' instrumental learning and action selection. The inability to accurately predict the timing of future outcomes diminishes an individual's perceived controllability of the environment, which has also been implicated in psychiatric disorders, such as anxiety (Bishop & Gagne, 2018).

Hidden-state inference models capture a similar idea as the microstimulus model at a different level of analysis (Starkweather, Babayan, Uchida, & Gershman, 2017). Often, the true state of the world is unknown or hidden and must be inferred from observations. This inference process is in part driven by prediction errors (Rouhani et al., 2020), and by extension is more difficult in volatile environments. As a result, ELU individuals may infer fewer states in the world (or, analogously, more states in an environment where negative prediction errors predominate) and group their experiences accordingly as a result of this early volatility. We have previously shown that this assumption of reduced sensitivity with a hidden-state inference model can produce reduced exploration in a foraging task (Harhen & Bornstein, 2023), a behavior found in ELU populations (Lloyd, McKay, & Furl, 2022), and may also explain why individuals who experience ELU are at higher risk of developing substance use disorders and relapsing following treatment (Harhen, Baram, Yassa, and Bornstein, 2021).

Our model is predicated on the assumption that prediction error learning can serve as a mechanism of environmental adaptation across multiple timescales—within a task and across development. Embodying an extreme form of sensitive period, adulthood is conceptualized as a period in which learning has altogether ceased. Future work could examine the effect of more realistic, relaxed constraints on learning in adulthood—in which developmental experience lays the groundwork for the architecture of neural systems which later adulthood experience can modify and reorganize (Galván, 2010; Karmiloff-Smith, 1994). Under this scenario, the prior biases instilled by the developmental environment should have their greatest influence in few shot or one shot learning experiences. When current experience underdetermines

what an agent should expect or do, past experience should largely influence the conclusion an agent reaches, with early life experience having a particularly privileged role (Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010). Such inductive biases facilitate learning in environments aligned with these biases and frustrate it in misaligned environments. If the influence of the developmental environment on expectations and choice is greatest in environments in which the agent has limited experience, this has implications for when symptoms for disorders like anxiety and substance use disorder should worsen (Bornstein & Pickard, 2020; Sharp, Miller, Dolan, & Eldar, 2020).

Our results highlight the key role time plays in shaping reinforcement learning and consequently its impact on behaviors associated with mental illness. The model's ability to produce varied phenotypes from the same computations suggests that the model's implications extend beyond anhedonia. Potentially, it provides a common origin for a number of psychiatric disorders, offering a potential explanation for high comorbidity rates (Jacobi et al., 2004; Kessler, Chiu, Demler, Merikangas, & Walters, 2005; Krueger, Chentsova-Dutton, Markon, Goldberg, & Ormel, 2003). Further research is needed to empirically test the model's behavioral predictions, namely, for ELU's impact on interval timing, and interval timing's relationship with psychiatric disorders. Finally, our results offer a demonstration of the value of computational modeling to understand the development of psychopathology. By drawing on a reinforcement learning framework, we can formalize the changing relationship between the agent and their environment across development, produce testable predictions of how the environment shapes the latent computations underlying clinically relevant behaviors, like learning, and propose mechanistic links between altered computations and the later emergence of psychiatric symptoms.

Acknowledgments

This work was supported by NIH P50MH096889 (PI: Tallie Z. Baram) and a NARSAD Young Investigator Award from the Brain and Behavior Research Foundation to AMB. NCH was supported by a National Defense Science and Engineering Graduate fellowship. The authors thank Elliot Ludvig for providing the microstimulus model code, and Tallie Z. Baram, Meredith Doherty, Steven J. Granger, Steven V. Mahler, Sophia C. Levis, and Michael A. Yassa for helpful conversations.

References

- Balsam, P. D., Drew, M. R., & Gallistel, C. R. (2010). Time and associative learning. *Comparative Cognition and Behavior Reviews*, 5, 1–22.
- Baram, T. Z., Davis, E. P., Obenaus, A., Sandman, C. A., Small, S. L., Solodkin, A., & Stern, H. (2012). Fragmentation and unpredictability of early-life experience in mental disorders. *American Journal of Psychiatry*, 169(9), 907–915.
- Belsky, J., & Fearon, R. M. P. (2002). Early attachment security, subsequent maternal sensitivity, and later child development: Does continuity in development depend upon continuity of caregiving? *Attachment & Human Development*, 4(3), 361–387.

- Birn, R. M., Roeber, B. J., & Pollak, S. D. (2017). Early childhood stress exposure, reward pathways, and adult decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(51), 13549–13554.
- Birnie, M. T., Kooiker, C. L., Short, A. K., Bolton, J. L., Chen, Y., & Baram, T. Z. (2020). Plasticity of the reward circuitry after early-life adversity: Mechanisms and significance. *Biological Psychiatry*, *87*(10), 875–884.
- Bishop, S. J., & Gagne, C. (2018). Anxiety, depression, and decision making: A computational perspective. *Annual Review of Neuroscience*, *41*, 371–388.
- Boecker, R., Holz, N. E., Buchmann, A. F., Blomeyer, D., Plichta, M. M., Wolf, I., Baumeister, S., Meyer-Lindenberg, A., Banaschewski, T., Brandeis, D., & Laucht, M. (2014). Impact of early life adversity on reward processing in young adults: EEG-fMRI results from a prospective study over 25 years. *PLoS One*, *9*(8), e104185.
- Bolton, J. L., Molet, J., Regev, L., Chen, Y., Rismanchi, N., Haddad, E., Yang, D. Z., Obenaus, A., & Baram, T. Z. (2018). Anhedonia following early-life adversity involves aberrant interaction of reward and anxiety circuits and is reversed by partial silencing of amygdala corticotropin-releasing hormone gene. *Biological Psychiatry*, *83*(2), 137–147.
- Bornstein, A. M., Aly, M., Feng, S. F., Turk-Browne, N. B., Norman, K. A., & Cohen, J. D. (2023). Associative memory retrieval modulates upcoming perceptual decisions. *Cognitive, Affective, Behavioral Neuroscience*, *23*(3), 645–665.
- Bornstein, A. M., & Pickard, H. (2020). “Chasing the first high”: Memory sampling in drug choice. *Neuropsychopharmacology*, *45*(6), 907–915.
- Cohen, J. (1992). Statistical power analysis. *Current Directions in Psychological Science*, *1*(3), 98–101.
- Cruz, F. C., Quadros, I. M., Planeta, C. d. S., & Miczek, K. A. (2008). Maternal separation stress in male mice: Long-term increases in alcohol intake. *Psychopharmacology*, *201*(3), 459–468.
- Day, C. V., Gatt, J. M., Etkin, A., DeBattista, C., Schatzberg, A. F., & Williams, L. M. (2015). Cognitive and emotional biomarkers of melancholic depression: An iSPOT-D report. *Journal of Affective Disorders*, *176*, 141–150.
- Dennison, M. J., Rosen, M. L., Sambrook, K. A., Jenness, J. L., Sheridan, M. A., & McLaughlin, K. A. (2019). Differential associations of distinct forms of childhood adversity with neurobehavioral measures of reward processing: A developmental pathway to depression. *Child Development*, *90*(1), e96–e113.
- Dillon, D. G., Holmes, A. J., Birk, J. L., Brooks, N., Lyons-Ruth, K., & Pizzagalli, D. A. (2009). Childhood adversity is associated with left basal ganglia dysfunction during reward anticipation in adulthood. *Biological Psychiatry*, *66*(3), 206–213.
- Dillon, D. G., & Pizzagalli, D. A. (2018). Mechanisms of memory disruption in depression. *Trends in Neurosciences*, *41*(3), 137–149.
- Dubal, S., Pierson, A., & Jouvent, R. (2000). Focused attention in anhedonia: A P3 study. *Psychophysiology*, *37*(5), 711–714.
- Efrati, A., & Gutfreund, Y. (2011). Early life exposure to noise alters the representation of auditory localization cues in the auditory space map of the barn owl. *Journal of Neurophysiology*, *105*(5), 2522–2535.
- Eldar, E., Rutledge, R. B., Dolan, R. J., & Niv, Y. (2016). Mood as representation of momentum. *Trends in Cognitive Sciences*, *20*(1), 15–24.
- Felitti, V. J. (2002). The relation between adverse childhood experiences and adult health: Turning gold into lead. *Permanente Journal*, *6*(1), 44–47.
- Frank, M. J. (2004). *Dynamic dopamine modulation of striato-cortical circuits in cognition: Converging neuropsychological, psychopharmacological and computational studies*. PhD thesis, University of Colorado at Boulder, Ann Arbor.
- Galván, A. (2010). Neural plasticity of development and learning. *Human Brain Mapping*, *31*(6), 879–890.
- Glynn, L. M., Stern, H. S., Howland, M. A., Risbrough, V. B., Baker, D. G., Nievergelt, C. M., Baram, T. Z., & Davis, E. P. (2019). Measuring novel antecedents of mental illness: The questionnaire of unpredictability in childhood. *Neuropsychopharmacology*, *44*(5), 876–882.

- Goff, B., Gee, D. G., Telzer, E. H., Humphreys, K. L., Gabard-Durnam, L., Flannery, J., & Tottenham, N. (2013). Reduced nucleus accumbens reactivity and adolescent depression following early-life stress. *Neuroscience*, 249, 129–138.
- Gollan, J. K., Pane, H. T., McCloskey, M. S., & Coccaro, E. F. (2008). Identifying differences in biased affective information processing in major depression. *Psychiatry Research*, 159(1–2), 18–24.
- Granger, S. J., Glynn, L. M., Sandman, C. A., Small, S. L., Obenaus, A., Keator, D. B., Baram, T. Z., Stern, H., Yassa, M. A., & Davis, E. P. (2021). Aberrant maturation of the uncinate fasciculus follows exposure to unpredictable patterns of maternal signals. *Journal of Neuroscience*, 41(6), 1242–1250.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences*, 14(8), 357–364.
- Hane, A. A., Henderson, H. A., Reeb-Sutherland, B. C., & Fox, N. A. (2010). Ordinary variations in human maternal caregiving in infancy and biobehavioral development in early childhood: A follow-up study. *Developmental Psychobiology*, 52(6), 558–567.
- Hanson, J. L., Albert, D., Iselin, A.-M. R., Carré, J. M., Dodge, K. A., & Hariri, A. R. (2016). Cumulative stress in childhood is associated with blunted reward-related brain activity in adulthood. *Social Cognitive and Affective Neuroscience*, 11(3), 405–412.
- Hanson, J. L., Williams, A. V., Bangasser, D. A., & Peña, C. J. (2021). Impact of early life stress on reward circuit function and regulation. *Frontiers in Psychiatry*, 12, 1799.
- Harhen, N., Baram, T. Z., Yassa, M. A., & Bornstein, A. M. (2021). Formalizing the relationship between early life adversity and addiction vulnerability: The role of memory sampling. *Biological Psychiatry*, 89(9), S189.
- Harhen, N., & Bornstein, A. M. (2023). Overharvesting in human patch foraging reflects rational structure learning and adaptive planning. *Proceedings of the National Academy of Sciences of the United States of America*, 120(13), e2216524120.
- Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*, 34(13), 4692–4707.
- Jacobi, F., Wittchen, H.-U., Holting, C., Höfler, M., Pfister, H., Müller, N., & Lieb, R. (2004). Prevalence, comorbidity and correlates of mental disorders in the general population: Results from the German Health Interview and Examination Survey (GHS). *Psychological Medicine*, 34(4), 597–611.
- Jin, D. Z., Fujii, N., & Graybiel, A. M. (2009). Neural representation of time in cortico-basal ganglia circuits. *Proceedings of the National Academy of Sciences of the United States of America*, 106(45), 19156–19161.
- Karmiloff-Smith, B. A. (1994). Beyond modularity: A developmental perspective on cognitive science. *European Journal of Disorders of Communication*, 29(1), 95–105.
- Kessler, R. C., Chiu, W. T., Demler, O., Merikangas, K. R., & Walters, E. E. (2005). Prevalence, severity, and comorbidity of 12-month DSM-IV disorders in the national comorbidity survey replication. *Archives of General Psychiatry*, 62(6), 617–627.
- Kim, G., Lewis-Peacock, J. A., Norman, K. A., & Turk-Browne, N. B. (2014). Pruning of memories by context-based prediction error. *Proceedings of the National Academy of Sciences of the United States of America*, 111(24), 8997–9002.
- Kosten, T. A., Miserendino, M. J., & Kehoe, P. (2000). Enhanced acquisition of cocaine self-administration in adult rats with neonatal isolation stress experience. *Brain Research*, 875(1–2), 44–50.
- Kosten, T. A., Sanchez, H., Zhang, X. Y., & Kehoe, P. (2004). Neonatal isolation enhances acquisition of cocaine self-administration and food responding in female rats. *Behavioural Brain Research*, 151(1–2), 137–149.
- Kosten, T. A., Zhang, X. Y., & Kehoe, P. (2006). Heightened cocaine and food self-administration in female rats with neonatal isolation experience. *Neuropsychopharmacology*, 31(1), 70–76.
- Krueger, R. F., Chentsova-Dutton, Y. E., Markon, K. E., Goldberg, D., & Ormel, J. (2003). A cross-cultural study of the structure of comorbidity among common psychopathological syndromes in the general health care setting. *Journal of Abnormal Psychology*, 112(3), 437–447.

- Lambert, H. K., King, K. M., Monahan, K. C., & McLaughlin, K. A. (2017). Differential associations of threat and deprivation with emotion regulation and cognitive control in adolescence. *Development and Psychopathology*, 29(3), 929–940.
- Li, Y., Fitzpatrick, D., & White, L. E. (2006). The development of direction selectivity in ferret visual cortex requires early visual experience. *Nature Neuroscience*, 9(5), 676–681.
- Lloyd, A., McKay, R. T., & Furl, N. (2022). Individuals with adverse childhood experiences explore less and underweight reward feedback. *Proceedings of the National Academy of Sciences of the United States of America*, 119(4). e2109373119.
- Ludvig, E. A., Sutton, R. S., & Kehoe, E. J. (2008). Stimulus representation and the timing of reward-prediction errors in models of the dopamine system. *Neural Computation*, 20(12), 3034–3054.
- Machlin, L., Miller, A. B., Snyder, J., McLaughlin, K. A., & Sheridan, M. A. (2019). Differential associations of deprivation and threat with cognitive control and fear conditioning in early childhood. *Frontiers in Behavioral Neuroscience*, 13, 80.
- McLaughlin, K. A., Sheridan, M. A., Humphreys, K. L., Belsky, J., & Ellis, B. J. (2021). The value of dimensional models of early experience: Thinking clearly about concepts and categories. *Perspectives on Psychological Science*, 16(6), 1463–1472.
- Mehta, M. A., Gore-Langton, E., Golembo, N., Colvert, E., Williams, S. C. R., & Sonuga-Barke, E. (2010). Hyporesponsive reward anticipation in the basal ganglia following severe institutional deprivation early in life. *Journal of Cognitive Neuroscience*, 22(10), 2316–2325.
- Miller, A. B., Machlin, L., McLaughlin, K. A., & Sheridan, M. A. (2021). Deprivation and psychopathology in the fragile families study: A 15-year longitudinal investigation. *Journal of Child Psychology and Psychiatry*, 62(4), 382–391.
- Molet, J., Maras, P. M., Kinney-Lang, E., Harris, N. G., Rashid, F., Ivy, A. S., Solodkin, A., Obenaus, A., & Baram, T. Z. (2016). MRI uncovers disrupted hippocampal microstructure that underlies memory impairments after early-life adversity. *Hippocampus*, 26(12), 1618–1632.
- NICHD Early Care Research Network. (2006). Infant–mother attachment classification: Risk and protection in relation to changing maternal caregiving quality. *Developmental Psychology*, 42(1), 38–58.
- Paine, T. A., Brainard, S., Keppler, E., Poyle, R., Sai-Hardebeck, E., Schwob, V., & Tannous-Taylor, C. (2021). Juvenile stress increases cocaine-induced impulsivity in female rats. *Behavioural Brain Research*, 414, 113488.
- Pike, A. C., & Robinson, O. J. (2022). Reinforcement learning in patients with mood and anxiety disorders vs control individuals: A systematic review and meta-analysis. *JAMA Psychiatry*, 79(4), 313–322.
- Rosen, M. L., Hagen, M. P., Lurie, L. A., Miles, Z. E., Sheridan, M. A., Meltzoff, A. N., & McLaughlin, K. A. (2020). Cognitive stimulation as a mechanism linking socioeconomic status with executive function: A longitudinal investigation. *Child Development*, 91(4), e762–e779.
- Rouhani, N., & Niv, Y. (2019). Depressive symptoms bias the prediction-error enhancement of memory towards negative events in reinforcement learning. *Psychopharmacology*, 236(8), 2425–2435.
- Rouhani, N., Norman, K. A., Niv, Y., & Bornstein, A. M. (2020). Reward prediction errors create event boundaries in memory. *Cognition*, 203, 104269.
- Schmidt, M. V. (2011). Animal models for depression and the mismatch hypothesis of disease. *Psychoneuroendocrinology*, 36(3), 330–338.
- Sharot, T. (2011). The optimism bias. *Current Biology*, 21(23), R941–R945.
- Sharp, P. B., Miller, G. A., Dolan, R. J., & Eldar, E. (2020). Towards formal models of psychopathological traits that explain symptom trajectories. *BMC Medicine*, 18(1), 264.
- Sheridan, M. A., Peverill, M., Finn, A. S., & McLaughlin, K. A. (2017). Dimensions of childhood adversity have distinct associations with neural systems underlying executive functioning. *Development and Psychopathology*, 29(5), 1777–1794.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24, 1193–1216.
- Sinclair, A. H., & Barense, M. D. (2018). Surprise and destabilize: Prediction error influences episodic memory reconsolidation. *Learning & Memory*, 25(8), 369–381.

- Smith, K. E., & Pollak, S. D. (2021). Rethinking concepts and categories for understanding the neurodevelopmental effects of childhood adversity. *Perspectives on Psychological Science*, 16(1), 67–93.
- Spadoni, A. D., Vinograd, M., Cuccurazzu, B., Torres, K., Glynn, L. M., Davis, E. P., Baram, T. Z., Baker, D. G., Nievergelt, C. M., & Risbrough, V. B. (2022). Contribution of early-life unpredictability to neuropsychiatric symptom patterns in adulthood. *Depression and Anxiety*, 39(10–11), 706–717.
- Sroufe, L. A. (2005). Attachment and development: A prospective, longitudinal study from birth to adulthood. *Attachment & Human Development*, 7(4), 349–367.
- Starkweather, C. K., Babayan, B. M., Uchida, N., & Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. *Nature Neuroscience*, 20(4), 581–589.
- Stuart, S. A., Hinchcliffe, J. K., & Robinson, E. S. J. (2019). Evidence that neuropsychological deficits following early life adversity may underlie vulnerability to depression. *Neuropsychopharmacology*, 44(9), 1623–1630.
- Sumner, J. A., Colich, N. L., Uddin, M., Armstrong, D., & McLaughlin, K. A. (2019). Early experiences of threat, but not deprivation, are associated with accelerated biological aging in children and adolescents. *Biological Psychiatry*, 85(3), 268–278.
- Sun, Y., Fang, J., Wan, Y., Su, P., & Tao, F. (2020). Association of early-life adversity with measures of accelerated biological aging among children in China. *JAMA Network Open*, 3(9), e2013588.
- Tanaka, S., Ribot, J., Imamura, K., & Tani, T. (2006). Orientation-restricted continuous visual exposure induces marked reorganization of orientation maps in early life. *Neuroimage*, 30(2), 462–477.
- Wakeford, A. G. P., Morin, E. L., Bramlett, S. N., Howell, B. R., McCormack, K. M., Meyer, J. S., Nader, M. A., Sanchez, M. M., & Howell, L. L. (2019). Effects of early life stress on cocaine self-administration in post-pubertal male and female rhesus macaques. *Psychopharmacology*, 236(9), 2785–2796.
- White, D. A., Myerson, J., & Hale, S. (1997). How cognitive is psychomotor slowing in depression? Evidence from a meta-analysis. *Neuropsychology, Development, and Cognition. Section B, Aging, Neuropsychology and Cognition*, 4(3), 166–174.
- Zhang, X. Y., Sanchez, H., Kehoe, P., & Kosten, T. A. (2005). Neonatal isolation enhances maintenance but not reinstatement of cocaine self-administration in adult male rats. *Psychopharmacology*, 177(4), 391–399.