

## Supplementary Materials

### Individuals with Adverse Childhood Experiences Explore Less and Underweight

#### Reward Feedback

Alex Lloyd, Ryan Thomas McKay & Nicholas Furl

*Proceedings of the National Academy of Sciences*

#### Contents

1. Instructions to participants
2. Full outline of the foraging task
3. Computational modelling
  - 3.1. Parameter recovery
  - 3.2. Model comparison
  - 3.3. Post-hoc simulation
4. Supplementary results
  - 4.1. Analyses controlling for gender
  - 4.2. Regression model statistics
  - 4.3. Analyses of ACE score as a continuous measure
  - 4.4. Analyses of participants' patch residency times
  - 4.5. Analyses controlling for counterbalance order effects
  - 4.6. Analyses of participants' reaction times
  - 4.7. Exploration behaviour in the high ACE subgroup
  - 4.8. Analyses by type of adversity

**1. Instructions to participants (note, bold text was also highlighted when instructions were provided to participants)**

*Instructions presented prior to the practice task*

Please read these instructions carefully before you begin.

Your aim is to collect as many apples as possible within the **time limit**. The more apples you collect, the larger your score at the end of this experiment and the bigger your prize will be.

- You can either **stay** to continue picking apples from the current tree or **leave** and find a new tree. If you leave and travel to a new tree, you have to wait a fixed amount of time. This time is fixed and has nothing to do with your internet connection or page loading.
- You will only need your keyboard for this task. You can either press 'S' to stay with the tree or 'L' to leave the tree and find a new one.
- You will not know how many apples are on a new tree until you stay and pick them, so it is a good idea to stay with each tree at least once before moving on.
- The number of apples left on a tree will decrease with time, meaning there will be fewer apples left on the tree to collect the longer you stay there. Apples do not grow back on each tree, so your job is to decide how long you want to spend at each tree.

**After seven minutes you will move into a completely new environment (think of it as a new orchard). This environment may be richer or poorer than the others. In some environments it may be better to stay with a tree for longer and in others it may be better to stay with a tree for less time.**

You will now begin a quick practice run of the study. Your practice score will not count towards your score in the main task. The task should take 15 minutes altogether (including the practice).

*Second set of instructions presented after practice*

You will now start the real study.

You will have seven minutes in each environment (or orchard) and the aim is to collect as many apples as possible as these count towards your total score. This should take no longer than 14 minutes to complete

Remember, it is a good idea to stay with each tree at least once to see how many apples are there. Apples do not grow back on each tree, so your job is to decide how long you want to spend at each tree. If you stay until there are no more apples left on the tree, you will not gain any more points.

Press the button below when you are ready.

## **2. Full Outline of the Foraging Task**


## Please read these instructions carefully before you begin.

Your aim is to collect as many apples as possible within the **time limit**. The more apples you collect, the larger your score at the end of this experiment and the bigger your prize will be.

- You can either **stay** to continue picking apples from the current tree or **leave** and find a new tree. If you leave and travel to a new tree, you have to wait a fixed amount of time. This time is fixed and has nothing to do with your internet connection or page loading.
- You will only need your keyboard for this task. You can either press 'S' to stay with the tree or 'L' to leave the tree and find a new one.
- You will not know how many apples are on a new tree until you stay and pick them, so it is a good idea to stay with each tree at least once before moving on.
- The number of apples left on a tree will decrease with time, meaning there will be fewer apples left on the tree to collect the longer you stay there. Apples do not grow back on each tree, so your job is to decide how long you want to spend at each tree.

**After seven minutes you will move into a completely new environment (think of it as a new orchard). This environment may be richer or poorer than the others. In some environments it may be better to stay with a tree for longer and in others it may be better to stay with a tree for less time.**

You will now begin a quick practice run of the study. Your practice score will not count towards your score in the main task. The task should take 15 minutes altogether (including the practice).




Practice Score: 10

You will now start the real study.

You will have seven minutes in each environment (or orchard) and the aim is to collect as many apples as possible as these count towards your total score. This should take no longer than 14 minutes to complete

Remember, it is a good idea to stay with each tree at least once to see how many apples are there. Apples do not grow back on each tree, so your job is to decide how long you want to spend at each tree. If you stay until there are no more apples left on the tree, you will not gain any more points.

Press the button below when you are ready.





You have now finished the task. Please take a moment to have a break.

In this section you will be asked some questions about experiences in your childhood. Some of these questions may be sensitive in nature.

All responses you give will be completely confidential and you will not be identified in any reports. You are free to leave out any questions you do not wish to answer.

When you are ready, please click the 'Next' button below.




*Supplementary Figure 1: Order with which participants completed the foraging task. The screens detail (in order): the initial instructions (see 1. Instructions to participant), the practice environment, final instructions, first environment (note, the order with which the colours were presented was counterbalanced), interim screen signalling the change in environment, the second environment and finally a break screen notifying participants about the upcoming ACE questionnaire.*

### **3. Computational Modelling**

#### **3.1. Parameter Recovery**

To ensure the model could estimate the correct value of the three free parameters, we simulated data for the rich and poor environments used in the behavioural study. For each free parameter, we simulated data across a range of values ( $\alpha$  {0,1} in increments of 0.1;  $c$  {-1, 3} in increments of 0.2 and  $\beta$  {1, 5} in increments of 0.4) and ran the computational model over these data. For our simulated and behavioural data, we estimated the parameters for each individual using the `optim` function<sup>1</sup> for R v.3.5.1. This function uses Maximum Likelihood Estimation to optimise the negative log likelihood of participants' stay/leave decisions produced by the SoftMax function (Equation 2). We found that the parameters were recoverable, meaning that there was a strong correlation between the parameters hard coded into the simulation and those estimated by the model (alpha:  $r(999) = .94$ ,  $p < .001$ , 95% CI [0.93, 0.94], beta:  $r(999) = .82$ ,  $p < .001$ , 95% CI [0.80, 0.84],  $c$ :  $r(999) = .31$ ,  $p < .001$ , 95% CI [0.25, 0.36]). This suggests that the model was able to estimate the correct value of the free parameters, though we found that  $c$  was less well recovered. Notably, we did not observe correlations between the parameters<sup>2</sup>. For correlations between the simulated and recovered parameters, see Supplementary Figure 2.



Supplementary Figure 2: Heatmap demonstrating the correlations between simulated (i.e., variables with the 'hardcode' prefix) and recovered parameters. Figures in each grid square indicate the Pearson's  $r$  correlation coefficient.

### 3.2. Model Comparison Analysis

We compared two models based on reinforcement learning to examine which of these presented a better explanation to participants' behaviour. The first model had a single learning rate parameter, which measured how much participants weighted reward feedback during their decision-making. The formalisation for this model is detailed in the main manuscript (see Equations 1 and 2). We compared this single learning rate model to a model which had separate learning rates for positive ( $\alpha_+$ ) and negative ( $\alpha_-$ ) prediction errors. In this model, we calculate whether reward feedback on each trial is better-than-expected (i.e., a positive prediction error) or poorer-than-expected (i.e., a negative prediction error), which is expressed as:

$$\delta_i \leftarrow s_i/\tau_i - \rho_i$$

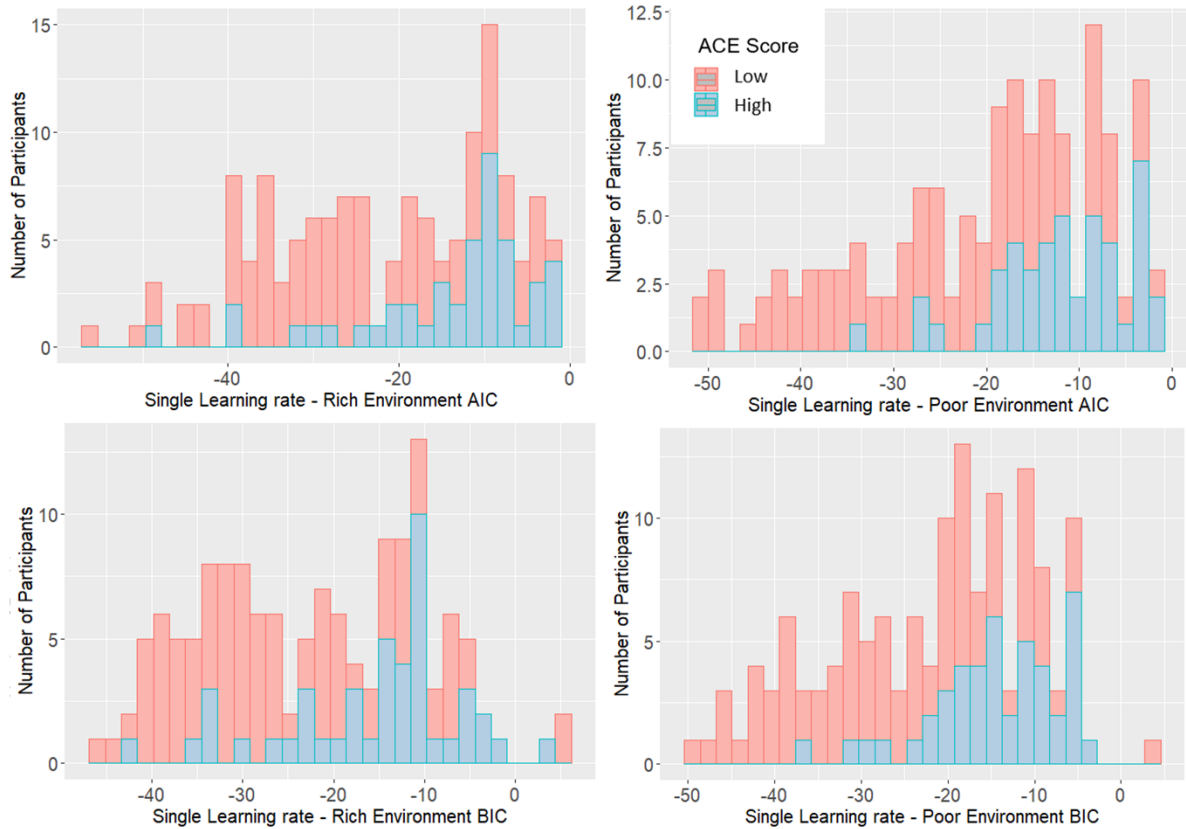
When  $\delta > 0$ , participants' expectation of the average reward rate ( $\rho$ ), as formulated in Equation 1, is updated according to the learning rate for positive prediction errors ( $\alpha_+$ ). When  $\delta < 0$ , participants' estimate of the average reward rate, as formulated in Equation 1, is updated according to the learning rate for negative prediction errors ( $\alpha_-$ ). Participants' estimate of the average reward rate is then entered into the SoftMax equation detailed in Equation 2 in the main manuscript.

We compared the models using two indices of model fit: the Bayesian Information Criterion (BIC) and Akaike's Information Criterion (AIC), where lower values indicate the model provides a better explanation of participants' data<sup>2</sup>. Comparison of the AIC and BIC for both models demonstrated that the single learning rate model had lower AIC and BIC values than the dual learning rate model in both the rich and poor foraging environments (see Supplementary Table 1), which was true for participants regardless of ACE exposure (see Supplementary Figure 3). These results suggest that a model with a single learning rate for both positive and negative prediction errors was a better fit to the data across both environments. As such, analyses in the main text were conducted using this model.

*Supplementary Table 1: Average of the indices of model fits from participants' data for the single and dual learning rate models, in each environment.*

Model	Rich Environment		Poor Environment	
	AIC	BIC	AIC	BIC
Dual learning rate model	-26.35	-14.61	-27.82	-17.97
Single learning rate model	-47.47	-39.49	-48.39	-40.86



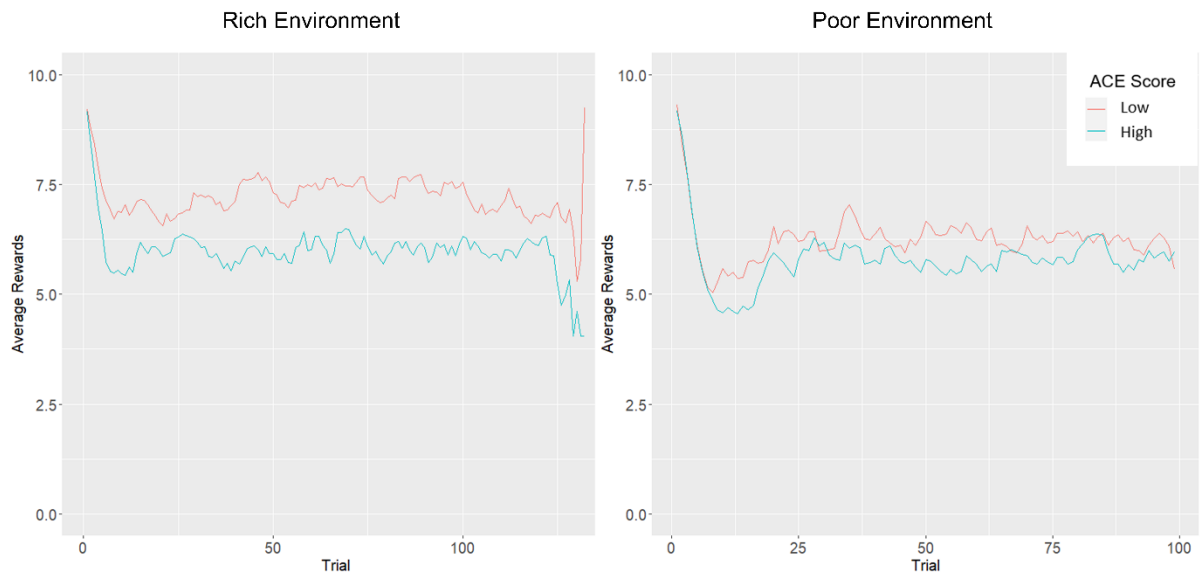


*Supplementary Figure 3: Plots demonstrating the relative fit of the single learning rate model compared to the dual learning rate model. Values below zero indicate that the single learning rate model was a better fit to participants' data whereas values above zero indicate the dual learning rate model was a better fit to the data.*

### 3.3. Post-Hoc Simulation

We used post-hoc simulation to demonstrate that we could recapitulate the main effect of ACE exposure reported in the Results section on participants' leaving threshold. We conducted simulations using the mean parameter from the single learning rate model for each ACE group in each foraging environment. A total of 200 agents were simulated, 100 using the mean parameter estimates from the high ACE group and 100 using the mean parameter estimates from the low ACE group. The plots in Supplementary Figure 4 descriptively indicate that our model was able to replicate the effect of ACE exposure on participants' leaving threshold (plotted in Figure 1 in the main manuscript), as we found the simulated 'high ACE' group to utilise a lower leaving threshold than the simulated 'low ACE group'. This group difference was confirmed statistically ( $F_{(1,198)} = 297.86, p < .001$ ). We also found that simulated

agents utilised a higher leaving threshold in the rich environment relative to the poor environment ( $F_{(1,198)} = 52.57, p < .001$ ), consistent with optimal foraging theory and participants' leaving threshold as reported in the Results section.



*Supplementary Figure 4: Plots demonstrating the average reward experienced on each trial for the simulated foraging agents in the rich environment (left) and poor environment (right). Local troughs indicate the simulated agents' leaving thresholds.*

## 4. Supplementary Results

### 4.1. Analyses Controlling for Gender

When controlling for gender in an ANCOVA of participants' leaving threshold, we preserved all the effects reported in the Results of the main manuscript, including the two main effects and null interaction effect. We found a main effect of ACE exposure ( $F_{(1,134)} = 4.84, p = .030, \eta^2 = .03$ ), and a main effect of environment type ( $F_{(1,134)} = 20.96, p < .001, \eta^2 = .002$ ). In addition, there was no interaction between ACE exposure and environment ( $F_{(1,134)} = 0.57, p = .452, \eta^2 < .001$ ).

When including gender as a covariate in an analysis of participants' learning rate, we also preserved the main effects and null interaction effect observed in the main manuscript. We found a main effect of ACE exposure ( $F_{(1,134)} = 9.26, p = .003, \eta^2 = .05$ ) and environment type ( $F_{(1,134)} = 5.13, p = .025, \eta^2 = .006$ ). Moreover, we did not observe an interaction between ACE exposure and environment ( $F_{(1,134)} = 0.36, p = .546, \eta^2 < .001$ ).

### 4.2. Regression Model Statistics

When each of the DOSEPERT subscales were entered into separate regression models, neither ACE score nor age predicted subscales of the DOSPERT. With the exception of the social risk subscale, we found that being male predicted the subscales of the DOSPERT (see Supplementary Table 2 for full model statistics).

Supplementary Table 2: Table demonstrating regression model statistics for the subscales of the DOSPERT. Bold cells indicate results that were statistically significant at the .05 level.

<i>Subscale</i>	<i>ACE Score</i>		<i>Age</i>		<i>Gender</i>		<i>Model Statistics</i>		
	$\beta$	$p$	$\beta$	$p$	$\beta$	$p$	$F$	$R^2$	$p$
<i>Social</i>	-0.09	.299	-0.12	.158	0.04	.631	1.17	0.03	.325
<i>Ethical</i>	0.10	.247	-0.06	.515	<b>0.19</b>	<b>.033</b>	1.73	0.04	.163
<i>Financial</i>	-0.13	.120	-0.11	.183	<b>0.25</b>	<b>.004</b>	<b>5.84</b>	<b>.001</b>	<b>0.11</b>

<i>Health and Safety</i>	0.15	.093	-0.10	.251	<b>0.23</b>	<b>.009</b>	<b>2.98</b>	<b>.033</b>	<b>0.06</b>
<i>Recreational</i>	-0.05	.588	-0.09	.272	<b>0.18</b>	<b>.044</b>	2.30	0.05	.080

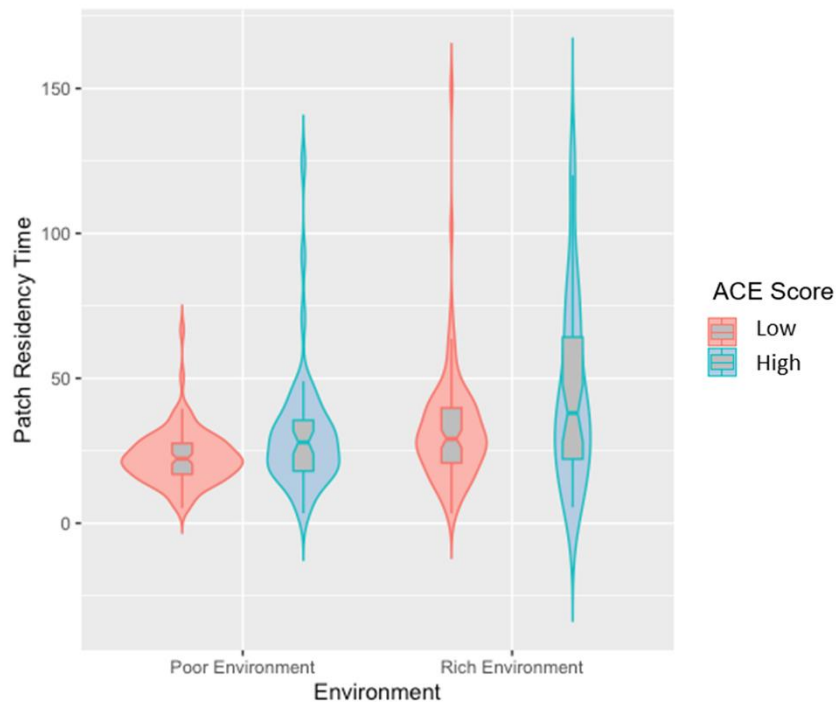
### 4.3. Analyses of ACE Score as a Continuous Measure

We found similar patterns in our data when nonparametric correlations were run using ACE scale as a continuous measure. Specifically, we found that there were significant negative associations between ACE scores and learning rate in the rich environment  $r(144) = -0.15, p = .033$  and poor environment  $r(144) = -0.20, p = .007$ . With regard to leaving thresholds, we found a significant negative association between ACE score and leaving threshold in the poor quality environment  $r(144) = -0.15, p = .033$ . However, the relationship between ACE score and leaving threshold in the rich environment was not significant  $r(144) = -0.09, p = .144$ .

### 4.4. Analyses of Participants' Patch Residency Times

One possibility is that the main effects of ACE group on participants' leaving threshold reported in the Results section occurred because participants in the high ACE group remained for the same or similar amount of time in patches in both environments, producing an illusory effect that they were adjusting their leaving thresholds between these two environments. To exclude this interpretation, we conducted analyses on participants' patch residency time measured in seconds (i.e., how long they remained in patches before exploring). Analysis of this variable replicated the findings we observed with participants' leaving thresholds. Specifically, participants adjusted their patch residency time as a function of the environment ( $F_{(1,136)} = 55.24, p < .001, \eta^2 = 0.07$ ), with participants remaining for longer in the richer environment relative to the poorer environment, which is consistent with optimal foraging theory<sup>3</sup>. We also found a main effect of ACE exposure ( $F_{(1,136)} = 11.78, p < .001, \eta^2 = 0.06$ ; see Supplementary Figure 5), consistent with our analyses of participants' leaving thresholds (see Figure 1 in the main manuscript). However, we did not find an interaction between environment

and ACE exposure ( $F_{(1,136)} = 1.31, p = .255, \eta^2 = 0.002$ ). Together, these findings suggest that participants adjusted the time they remained with patches in a similar manner to their leaving thresholds between the two foraging environments. Moreover, ACE exposed individuals stayed with patches for longer than individuals without these experiences.



*Supplementary Figure 5: Plot demonstrating the average patch residency time (y-axis) in the two foraging environments (x-axis). These results are comparable to our analyses of participants' leaving thresholds detailed in Figure 1 in the main manuscript.*

#### **4.5. Analyses Controlling for Counterbalance Order Effects**

We conducted additional analyses which suggested that our primary findings were not affected by order effects related to counterbalancing of the environments, as the effects reported in the Results section were preserved after accounting for the order with which participants experienced the two foraging environments. In these analyses we included the order with which participants experienced the rich and poor environments as an additional factor in the analyses of participants' leaving threshold and learning rates.

Examining participants' leaving thresholds, we did find a significant order\*environment interaction ( $F_{(1,136)} = 5.31, p = .023$ ). Despite this interaction, our results were consistent with

those reported in the main manuscript as we still find a main effect of ACE exposure ( $F_{(1,135)} = 4.39, p = .038, \eta^2 = 0.03$ ) and environment ( $F_{(1,135)} = 28.65, p < .001, \eta^2 = 0.03$ ), with the effect sizes for these differences being exactly the same as reported in our original analyses. We did not find a significant order\*environment\*ACE group interaction ( $F_{(1,135)} = 0.44, p = .510$ ), indicating these order effects did not account for the different leaving thresholds we observed between groups. As we had a priori predictions that leaving thresholds would be higher in the rich environment relative to the poor environment, we ran a one-tailed planned comparison between leaving thresholds in the rich and poor environments, accounting for counterbalance order. Results of this planned comparison demonstrated that leaving thresholds were significantly higher in the rich environment relative to the poor environment in the counterbalance condition where the rich environment was presented first ( $t(135) = 2.08, p = .039$ ) and in the counterbalance condition where the poor environment was presented first ( $t(135) = 5.51, p < .001$ ). Examining participants' learning rates, an ANOVA demonstrated that there was no interaction between counterbalance order\*environment ( $F_{(1,142)} = 0.01, p = .915$ ), though consistent with the results in the main manuscript we find a main effect of ACE exposure ( $F_{(1,142)} = 8.65, p = .004$ ) and environment ( $F_{(1,142)} = 0.01, p = .915$ ). These analyses demonstrate that the results reported in the original manuscript are not affected by order effects.

#### **4.6. Analyses of Participants' Reaction Times**

Analysing participants' reaction times, we found that participants in the high ACE group had slower reaction times ( $M_{rich} = 0.53, SD_{rich} = 0.24, M_{poor} = 0.51, SD_{poor} = 0.23$ ) compared to participants in the low ACE group ( $M_{rich} = 0.43, SD_{rich} = 0.17, M_{poor} = 0.43, SD_{poor} = 0.16; F_{(1,137)} = 8.53, p = .004, \eta^2 = 0.05$ ). However, we found no main effect of environment on RT ( $F_{(1,137)} = 0.39, p = .535$ ) nor an interaction between ACE group and environment ( $F_{(1,137)} = 0.84, p = .361$ ).

#### **4.7. Exploration Behaviour in the High ACE Group Only**

To exclude the possibility that participants in the high ACE group simply did not change their leaving thresholds much between the two environments, we ran independent *t*-tests on the leaving threshold and patch residency time in the high ACE group. These analyses demonstrated that participants in the high ACE group did modulate their leaving thresholds as a function of the environment, as these participants explored more in the rich environment ( $M = 5.00$ ,  $SD = 2.46$ ) than the poor environment ( $M = 4.26$ ,  $SD = 2.16$ )  $t(44) = 3.35$ ,  $p = .002$ , Cohen's  $d = 0.50$ . In addition, we also found that the high ACE group had significantly longer patch residency times in the rich environment relative to the poorer environment ( $t(44) = 4.59$ ,  $p < .001$ , Cohen's  $d = 0.68$ ), which is consistent with optimal foraging theory<sup>3</sup>.

#### **4.8. Analyses by Type of Adversity**

It has been suggested by some researchers (e.g.,<sup>4</sup>) that behavioural differences following adversity are driven by the type of experience. These experiences can be separated into the following categories: threatening events (including physical and sexual abuse), neglect, and family adversity (e.g., parental imprisonment; see Table 1 in the main manuscript;<sup>4-6</sup>). As some participants experienced more than one type of adversity, we constructed three factors: One that dummy coded for the presence versus absence of threatening events, one that coded for the presence versus absence of neglect, and one that coded for the presence versus absence of family adversity. For each of our dependent variables (leaving thresholds and learning rate), we ran separate ANOVAs for each of the three adversity factors to examine whether behavioural differences were driven by a particular type of adverse experience. We also included environment type (rich or poor) as a second factor in these ANOVA analyses.

In our analysis of participants' leaving thresholds, we did not find evidence that a particular category of adversity explains the differences we observe in participants' leaving thresholds. We ran three ANOVAs which each had only one of the types of adversity

(threatening events, neglect, or family adversity) entered as the independent variable. We applied the Bonferroni correction to these three exploratory analyses. When experience of threatening events was entered as the independent variable, we did not find a main effect of threatening events on participants' leaving thresholds ( $F_{(1,137)} = 2.17$ ,  $p_{bonf} = .429$ ), nor an interaction between threatening events and environment ( $F_{(1,137)} = 0.03$ ,  $p_{bonf} = 1$ ) though in this model we did find a main effect of environment, consistent with the analyses presented in the main manuscript ( $F_{(1,137)} = 28.83$ ,  $p_{bonf} = .003$ ). When experience of neglect was entered as the independent variable, we did not find a main effect of neglect ( $F_{(1,137)} = 0.21$ ,  $p_{bonf} = 1$ ) nor an interaction between environment and neglect ( $F_{(1,137)} = 2.61$ ,  $p_{bonf} = .327$ ), though we did find a main effect of environment, consistent with our original analyses ( $F_{(1,137)} = 31.38$ ,  $p_{bonf} = .003$ ). Finally, when experience of family adversity was entered as the independent variable, we did not find a main effect of family adversity ( $F_{(1,137)} = 3.73$ ,  $p_{bonf} = .168$ ), nor an interaction between family adversity and environment ( $F_{(1,137)} = 1.82$ ,  $p_{bonf} = .540$ ), though we did find a main effect of environment ( $F_{(1,137)} = 28.15$ ,  $p_{bonf} = .003$ ).

Although there were no effects of adversity type for leaving threshold, we found that group level differences in learning rate were driven by experiences of threatening events. We ran three ANOVAs examining the effect of each type of adversity on participants' learning rate. We Bonferroni-corrected for these three exploratory analyses. When experience of threatening events was entered as an independent variable, we found a main effect of threatening events on participants learning rate ( $F_{(1,137)} = 7.02$ ,  $p_{bonf} = .027$ ,  $\eta^2 = 0.05$ ) with participants exposed to these events having a lower learning rate ( $M_{rich} = 0.53$ ,  $SD_{rich} = 0.24$ ,  $M_{poor} = 0.59$ ,  $SD_{poor} = 0.17$ ) compared to participants who were not exposed to these experiences ( $M_{rich} = 0.62$ ,  $SD_{rich} = 0.14$ ,  $M_{poor} = 0.64$ ,  $SD_{poor} = 0.17$ ). We also found a main effect of environment ( $F_{(1,137)} = 9.31$ ,  $p_{bonf} = .009$ ,  $\eta^2 = 0.01$ ), though we did not find an interaction between environment and threatening events ( $F_{(1,137)} = 2.27$ ,  $p_{bonf} = .405$ ). When experience of neglect was entered as an independent variable, we did not find a main effect of neglect



( $F_{(1,137)} = 1.53$ ,  $p_{bonf} = .654$ ) nor an interaction between environment and neglect ( $F_{(1,137)} =$ ,  $p_{bonf} = .327$ ), though we did find a main effect of environment ( $F_{(1,137)} = 31.38$ ,  $p_{bonf} = .003$ ). Finally, when experience of family adversity was entered as an independent variable, we did not find a main effect of family adversity ( $F_{(1,137)} = 4.29$ ,  $p_{bonf} = .120$ ). We found a main effect of environment ( $F_{(1,137)} = 9.38$ ,  $p = .003$ ,  $p_{bonf} = .009$ ), though we did not find an interaction between environment and family adversity ( $F_{(1,137)} = 0.03$ ,  $p_{bonf} = 1$ ). These findings may suggest that differences in learning rate were driven by threatening experiences.

## Supplementary References

1. Nash, J. C. On Best Practice Optimization Methods in *R. J. Stat. Soft.* **60**, (2014).
2. Daw, N. D. Trial-by-trial data analysis using computational models. *Decision making, affect, and learning: Attention and performance XXIII* **23**, 26.
3. Gabay, A. S. & Apps, M. A. J. Foraging optimally in social neuroscience: computations and methodological considerations. *Social Cognitive and Affective Neuroscience* **16**, 782–794 (2021).
4. McLaughlin, K. A. & Sheridan, M. A. Beyond Cumulative Risk: A Dimensional Approach to Childhood Adversity. *Curr Dir Psychol Sci* **25**, 239–245 (2016).
5. Sheridan, M. A. & McLaughlin, K. A. Dimensions of early experience and neural development: deprivation and threat. *Trends in Cognitive Sciences* **18**, 580–585 (2014).
6. McLaughlin, K. A., Sheridan, M. A. & Lambert, H. K. Childhood Adversity and Neural Development: Deprivation and Threat as Distinct Dimensions of Early Experience. *Neurosci Biobehav Rev* **47**, 578–591 (2014).