

This is a repository copy of *A series of unfortunate events: Do those who catastrophize learn more after negative outcomes?*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/208867/>

Version: Published Version

Article:

Pike, Alex orcid.org/0000-0003-1972-5530, Harada-Laszlo, Mia, Talwar, Anahita et al. (1 more author) (2023) A series of unfortunate events: Do those who catastrophize learn more after negative outcomes? Mental Health Science. e49. ISSN: 2642-3588

<https://doi.org/10.1002/mhs2.49>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

RESEARCH ARTICLE

WILEY

A series of unfortunate events: Do those who catastrophize learn more after negative outcomes?

Mia Harada-Laszlo¹ | Anahita Talwar¹ | Oliver J. Robinson¹ | Alexandra C. Pike^{1,2} 

¹Anxiety Lab, Neuroscience and Mental Health Group, Institute of Cognitive Neuroscience, University College London, London, UK

²Department of Psychology, York Biomedical Research Institute, University of York, York, UK

Correspondence

Alexandra C. Pike, Department of Psychology, University of York, York YO10 5DD, UK.
Email: alex.pike@york.ac.uk

Oliver J. Robinson, Institute of Cognitive Neuroscience, University College London, London, WC1N 3AR.
Email: o.robinson@ucl.ac.uk

Funding information

Medical Research Council,
Grant/Award Number: MR/R020817/1

Abstract

Catastrophizing is a transdiagnostic construct that has been suggested to precipitate and maintain a multiplicity of psychiatric disorders, including anxiety, depression, post-traumatic stress disorder, and obsessive-compulsive disorder. However, the underlying cognitive mechanisms that result in catastrophizing are unknown. Relating reinforcement learning model parameters to catastrophizing may allow us to further understand the process of catastrophizing. Using a modified four-armed bandit task, we aimed to investigate the relationship between reinforcement learning parameters and self-report catastrophizing questionnaire scores to gain a mechanistic understanding of how catastrophizing may alter learning. We recruited 211 participants to complete a computerized four-armed bandit task and tested the fit of six reinforcement learning models on our data, including two novel models which both incorporated a scaling factor related to a *history of negative outcomes* variable. We investigated the relationship between self-report catastrophizing scores and free parameters from the overall best-fitting model, along with the best-fitting model to include *history*, using Pearson's correlations. Subsequently, we reassessed these relationships using multiple regression analyses to evaluate whether any observed relationships were altered when relevant IQ and mental health covariates were applied. Model-agnostic analyses indicated there were effects of outcome history on reaction time and accuracy, and that the effects on accuracy related to catastrophizing. The overall model of best fit was the Standard Rescorla–Wagner Model and the best-fitting model to include *history* was a model in which learning rate was scaled by history of negative outcome. We found no effect of catastrophizing on the scaling by history of negative outcome parameter ($r = 0.003$, $p = 0.679$), the learning rate parameter ($r = 0.026$, $p = 0.703$), or the inverse temperature parameter ($r = 0.086$, $p = 0.220$). We were unable to relate catastrophizing to any of the reinforcement learning parameters we investigated. This implies that catastrophizing is not straightforwardly linked to any changes to learning after a series of negative outcomes are received. Future research could incorporate further exploration of the space of models which include a *history* parameter.

Oliver J. Robinson and Alexandra C. Pike are joint senior authors.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Mental Health Science* published by Wiley Periodicals LLC.

KEYWORDS

catastrophizing, computational modeling, computational psychiatry, online, reinforcement learning, transdiagnostic

1 | INTRODUCTION

The term catastrophizing refers to the tendency of an individual to “[overestimate] the likelihood of a negative event, and also [believe] that the negative event will be catastrophic” (Pike et al., 2021), and has been a recurring concept across the psychiatric literature since first being coined in the 1960s (Ellis, 1962). In clinical populations, catastrophizing may present as overestimating the likelihood of a trauma reoccurring in those with post-traumatic stress disorder (Dunmore et al., 1999), or predicting excessively negative consequences in response to physical anxiety symptoms in those with panic disorder (Austin & Richards, 2001). However, it is largely accepted to be a phenomenon also present to varying degrees in nonclinical populations, in scenarios such as believing one's plane may crash in response to turbulence. It has been suggested that catastrophizing may be a transdiagnostic construct that contributes to both the onset (Jenness et al., 2016; McLaughlin et al., 2014; Paré et al., 2019) and maintenance of a multiplicity of psychiatric conditions (Gellatly & Beck, 2016), with the content of the “catastrophe” differing between disorders, but the mental processes and manifested behaviors (such as safety-seeking or an inability to reappraise a situation), remaining largely consistent. Notably, the transdiagnostic relevance of catastrophizing is well supported by the widespread success of decatastrophizing cognitive therapy in treating a range of disorders (Bowers et al., 1997; Clark & Beck, 2011).

In our previous work developing a self-report measure for catastrophizing for use in a broader psychiatric context (Pike et al., 2021), we found that our catastrophizing measure exhibited a greater ability to predict psychiatric diagnoses and psychiatric medications over and above the common, clinically implemented Patient Health Questionnaire-9 (PHQ-9) (Kroenke et al., 2001) and Generalized Anxiety Disorder Assessment (GAD-7) (Spitzer et al., 2006) questionnaires, demonstrating that catastrophizing has the potential to be an important predictor of mental health susceptibility. Despite this, there has been little detailed research exploring the cognitive processes that initiate this cycle.

In this study, we proposed the application of a reinforcement learning framework to explore whether computational methods can be used to capture catastrophizing cognitions, by parameterizing the psychopathological alterations of learning that may lead to maladaptive behaviors downstream.

Reinforcement learning is a learning algorithm in which an individual uses the disparity between expected and obtained outcomes (the prediction error) to update internal associations between certain stimuli or actions and their outcomes. This, in turn, signals the need to adjust behavior accordingly, and this process ultimately optimizes decision-making to maximize reward and minimize loss long term (Sutton & Barto, 1998).

The reinforcement learning framework has been successfully employed to investigate mental health conditions that involve catastrophizing (Aylward et al., 2019; Dayan, 2009; Kumar et al., 2008; Pike & Robinson, 2022), but never to empirically describe catastrophizing as an isolated construct. These reinforcement learning studies often implicitly assume that changes in learning might cause changes in mood or other symptoms, thus implicating these learning changes as targets for intervention. This assumption is plausible: there is evidence that cognitive changes are predictive of response to antidepressants (Park et al., 2018), and cognitive biases predict later onset of depression and anxiety (Smith et al., 2018). If this assumption holds, using reinforcement learning parameters to define how learning processes are altered in those who catastrophize could drive a more informed approach to dismantling catastrophic cognition by identifying specific targets for interventions, not only for translational use across conditions but also as an early intervention to protect at-risk populations from the onset of these disorders. Moreover, these methods represent a pivotal step towards advancing the precision with which we are able to assess the effectiveness of cognitive and psychopharmacological intervention, by providing a framework to measure changes in these reinforcement learning processes over time (Lawson et al., 2021), without the subjectivity of self-report. However, it may be noted that there is yet to be concrete evidence for this assumption of a causal pathway, and therefore, it is plausible that the direction of causality may be the reverse. Future longitudinal studies may be necessary to disentangle these constructs.

Given that catastrophizing is characterized in the literature as an exaggerated perception of the seriousness of the consequences of a negative or ambiguous event, we considered it plausible that catastrophizing could manifest as overweighting recent (negative) outcomes when learning, particularly when these are frequent. Computationally, this can be captured by scaling the *learning rate* parameter in a reinforcement learning model by the proportion of previous negative outcomes within a given window. In other words, following a string of negative outcomes, we suggest catastrophizers will update their learnt value for a stimulus to a greater degree than noncatastrophizers.

1.1 | Study objectives and hypotheses

1.1.1 | Objective

To investigate the relationship between reinforcement learning parameters and self-report catastrophizing questionnaire scores. By modeling a modified four-armed bandit task, we aim to elucidate the mechanisms by which catastrophizing may alter learning.

To address this objective, our study tested six learning models and investigated the following predictions, which were largely dependent on which model proved to be the best fit for our data, as stated in our preregistration document (Harada-Laszlo et al., 2021).

H1: First, we hypothesized that models that included the *history of negative outcome* as a parameter would fit the data better than those that lack this variable. Furthermore, we predicted that the best-fitting model to include a *history* term would be one that scales *learning rate* by *history* rather than one in which *sensitivity* is scaled by *history*.

H2.1: In the case that the best-fitting model was one that scales *learning rate* by *history of negative outcome*, we predicted that there would be a positive correlation between the *scaling* parameter and self-report catastrophizing scores.

H2.2: In the case that the best-fitting model was one in which the *sensitivity* parameter was scaled by *history of negative outcome*, we predicted there would be a positive correlation between the *sensitivity* parameter and self-report catastrophizing scores.

2 | MATERIALS AND METHODS

Note that this study was preregistered ([10.17605/OSF.IO/VSN8B](https://osf.io/VSN8B)), and open data and code can be found at [10.17605/OSF.IO/3Y6UJ](https://osf.io/3Y6UJ).

2.1 | Participants and recruitment

We recruited 211 participants for this study (145 females, mean \pm SD age = 40.1 \pm 10.7; 66 males, mean \pm SD age = 43.6 \pm 10.0) via the Prolific platform (<https://app.prolific.co/>), as informed by a priori power analysis (Supporting Information S1: Section 1.1). We set inclusion and exclusion criteria via the Prolific platform, details of which can be found in Supporting Information S1: Section 1.2.

2.2 | Experimental procedure

Following recruitment, participants were directed to the Gorilla platform, where we obtained informed consent and presented all components of the experiment. Participants were randomly assigned to one of four versions of the modified four-armed bandit task (see Figure 1c), and subsequently directed to complete the forward digit span (Wechsler, 1955) (Supporting Information S1: Section 3) and one of three versions of the adapted ART task (Chierchia et al., 2019) (Supporting Information S1: Section 4) in an order randomized between participants. Next, participants answered brief questions

about any mental health diagnoses and medications used, before commencing the questionnaire battery (Supporting Information S1: Section 5), in which all six questionnaires were counterbalanced among participants. The experimental session lasted approximately 26 min (mean \pm SD session duration = 25.6 min \pm 7.5 min).

2.3 | Four-armed bandit task

To capture participant learning and decision-making, we devised a computerized probabilistic instrumental learning task, based on the framework of the “four-armed bandit task” (Daw et al., 2006), a common learning paradigm that allows for the presentation of rewards and punishments (Averbeck, 2015).

On each trial, participants were given the choice of four “bandits” with fluctuating associations with appetitive or aversive outcomes and were instructed to choose the bandit most likely to deliver a reward. Feedback was displayed immediately following bandit selection and points were exhibited on screen throughout the experimental session. Participants were given 5 s to make their selection, before the task progressed automatically onto the next trial (trials where no choice was made were censored from modeling). Bandits remained in a fixed position throughout the experimental session, and there was always one bandit that would deliver a reward on each trial, with the choice of any other bandit resulting in negative feedback.

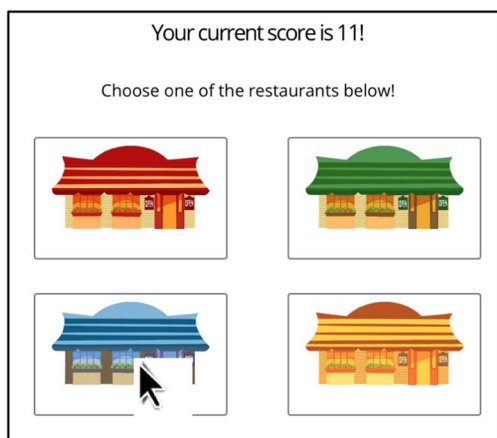
Our task consisted of 200 trials, and the reward probabilities for each bandit changed with each block of 25-trials, without the participant being informed. The order of these eight blocks was randomized between participants, and trials within each block were shuffled. For each block, there was one winning bandit with a pay-off between 50% and 75% (Figure 1b). Each of the four bandits was the winning bandit in two of the 25-trial blocks, to ensure continual learning of stimulus-outcome contingencies. For further task description see Supporting Information S1: Section 2.

2.4 | Additional tasks

We included two additional tasks and a battery of six short mental health related questionnaires, to measure the factors which we considered to be vital covariates for analysis.

We employed a computerized adaptation of the *forward digit span* as a proxy for working memory (Wechsler, 1955), as participants are required to remember and integrate recent outcome history for maximal performance in the reinforcement learning task (Collins et al., 2014). Moreover, we implemented a computerized *Abstract Reasoning Task* (Chierchia et al., 2019) as a proxy for fluid intelligence, which has been proposed as a determinant of choice and strategy in reinforcement learning (Schad et al., 2014). To keep our experiment within reasonable time limits, we developed a shortened 4-minute version of this task (Chierchia et al., 2019). Details of this procedure along with the analysis of the psychometric properties of the

(a) BANDIT SELECTION SCREEN (5s)




REWARD


OR

PUNISHMENT

FEEDBACK SCREEN



(b)



(c)

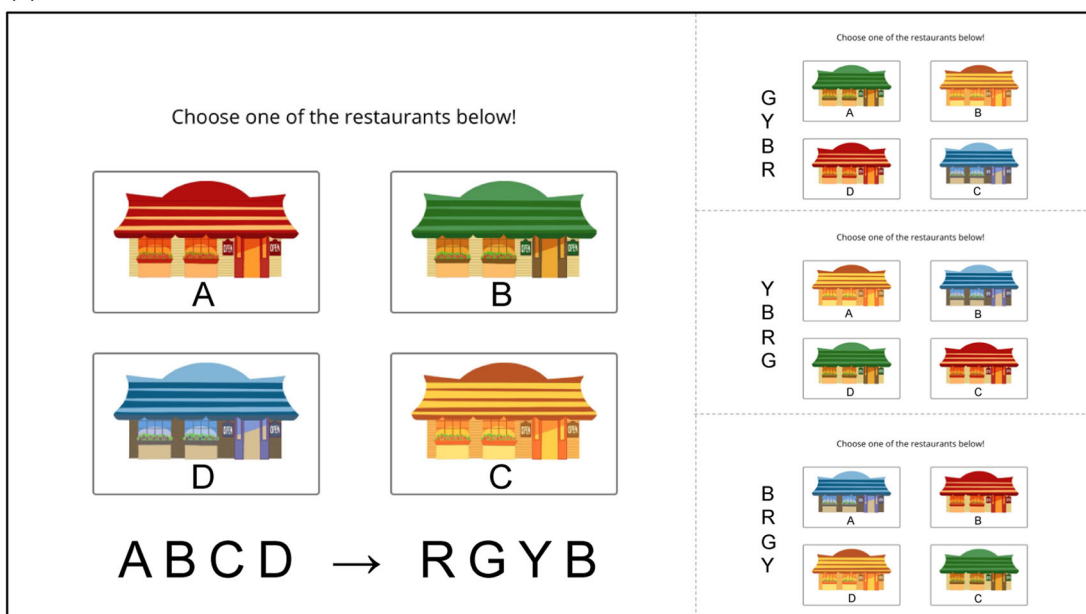


FIGURE 1 The modified four-armed bandit task. We constructed a narrative in which bandits were represented by images of restaurants, and the fluctuating probabilities of different outcomes was communicated via the instruction that chefs could change randomly at any time, and that even good chefs would sometimes make a bad meal. (a) demonstrates trial structure, including the bandit selection screen and feedback screen as seen by participants. (b) illustrates the fluctuation of reward count per bandit across 200 trials, where reward probabilities shift every 25 trials, and each color represents a bandit. The y axis indicates the number of times each bandit is the rewarded/correct choice per 25-trial block. (c) illustrates the four variations of stimuli color randomization. Bandits A, B, C, and D remained in a fixed position at top left, top right, bottom right, and bottom left respectively. Each of the colors, Red (R), Green (G), Yellow (Y), and Blue (B) were used to represent each bandit for one in four participants, in a manner randomized between participants.

shortened adaptation can be found in Supporting Information S1: Section 4.

Our questionnaire battery comprised the Self-Report Catastrophizing Questionnaire to assess catastrophizing (Pike et al., 2021), the GAD-7 to assess anxiety (Spitzer et al., 2006), the Spielberger State-Trait Anxiety Inventory to measure trait anxiety (Spielberger, 1983), the PHQ-8 to measure depression (Kroenke et al., 2009), the Penn State Worry Questionnaire (PSWQ) to measure worry (Meyer et al., 1990), and our Covid-19 Impact Questionnaire to assess the extent to which an individual has been affected by the Covid-19 pandemic (Pike et al., 2023). We included these measures to allow us to assess whether any results were specific to catastrophizing, rather than frequently co-existing psychopathology. Further details of each questionnaire can be found in Supporting Information S1: Section 5.

2.5 | Statistical analysis

2.5.1 | Model-agnostic analyses

We analyzed participants' task performance to understand whether there were any relationships between catastrophizing scores and model-agnostic measures of behavior. First, we performed simple Pearson's correlations between catastrophizing score and "task score" (defined as choosing the bandit that was rewarded on that given trial), and also correlated catastrophizing scores with each participants' win-stay proportion (their tendency to repeat their choice after receiving a win outcome) and their lose-shift proportion (their tendency to choose a different stimulus after not receiving a win outcome).

To assess in more detail the effect of outcome history on participants' performance, we performed two mixed model analyses with a random intercept for each participant. Please note that these were not preregistered, as they do not relate to our primary hypotheses (relationships between model parameters and catastrophizing). First, we performed a logistic mixed model of lagged outcome on participants' subsequent choice to switch (choose another stimulus) or stay (repeat the same choice), and a mixed model of lagged outcome against reaction time (RT). We also repeated these analyses including only the significant lag terms and adding catastrophizing scores.

2.5.2 | Learning models and model fitting

We then initially tested the fit of six learning models to our data, including the Standard Rescorla-Wagner Model (Rescorla, 1972), the Rescorla-Wagner model with a sensitivity term (Huys et al., 2013), the Vmax model (Huang et al., 2017) and the Pearce-Hall associability model (Pearce & Hall, 1980), along with two novel learning models designed to test our hypotheses, in which *learning rate* and *sensitivity* are respectively scaled by *history of negative outcome*. For model equations and descriptions, see Supporting Information S1: Section 6.

Models were written using Stan syntax and model fitting was conducted using CmdStanR. Our models were hierarchical in structure and as such, we derived model parameters from a group distribution, parameterized by weakly informative hyperparameters. We conducted posterior inference for our models using a Markov-Chain Monte-Carlo sampling scheme implemented in R v4.0.2 and CmdStanR v0.3.0. We used 4 chains with 2000 iterations each, of which we designated half for warmup. We employed the No-U-Turn Sampling algorithm (Hoffman & Gelman, 2014). We selected the best-fitting model as defined by the integrated Bayesian Information Criterion (iBIC) (Huys et al., 2012), which generates a score based on the likelihood (the probability that the model assigns to the actual choice made by the participant) and incorporates a penalty for increased model complexity.

2.6 | Data analysis

To assess relationships between catastrophizing and cognition, we performed Pearson's correlations between catastrophizing scores and all reinforcement learning parameters of the best-fitting model as well as the relevant *scaling* or *sensitivity* parameter in the best-fitting model that included *history*. Subsequently, we used multiple regression to investigate whether these relationships were affected by adding several potentially important covariates to the model, including working memory, fluid intelligence, age, and sex. In addition, we performed a second multiple regression to investigate the specificity of our findings to catastrophizing, by creating a complete model that incorporated all questionnaire scores, along with all covariates included in the first multiple regression.

3 | RESULTS

3.1 | Model-agnostic analyses

There was no correlation between catastrophizing scores and task scores ($r(223) = 0.098$, $p = 0.145$), nor was there any correlation between catastrophizing scores and lose-shift proportion ($r(233) = 0.094$, $p = 0.161$). There was a correlation between catastrophizing scores and win-stay proportion ($r(223) = 0.161$, $p = 0.015$), such that those with higher catastrophizing scores tended to repeat a choice after a win.

There was a significant effect of switch/stay choices for up to three trials previously (Table 1), such that after receiving a win outcome, individuals were more likely to "stay"—and this effect was true even if the win outcome was three trials ago. There was also a significant effect on reaction time from the two previous trials (Table 2): after a win outcome, participants responded more slowly for one trial, then more rapidly on the subsequent trial.

We also performed mixed model analyses to examine the interaction between z scored catastrophizing scores and the significant lag terms for both of the above models. There were

TABLE 1 Logistic mixed model of switch/stay on trial t (optimizer: bobyqa, nAGQ: 10).

Term	Odds ratio (CI)	p
Intercept	0.24 (0.20, 0.29)	<0.001***
Lag outcome ($t - 1$)	30.66 (28.66, 32.80)	<0.001***
Lag outcome ($t - 2$)	1.81 (1.81, 1.92)	<0.001***
Lag outcome ($t - 3$)	1.12 (1.05, 1.18)	<0.001***
Lag outcome ($t - 4$)	0.98 (0.92, 1.04)	0.429
Lag outcome ($t - 5$)	0.99 (0.93, 1.05)	0.676

Abbreviation: CI, confidence interval.

*** $p < 0.001$.**TABLE 2** Mixed model of reaction time on trial t .

Term	Estimate (CI)	p
Intercept	613.43 (577.09, 649.76)	<0.001***
Lag outcome ($t - 1$)	33.24 (20.51, 45.97)	<0.001***
Lag outcome ($t - 2$)	-13.38 (-26.44, -0.32)	<0.045*
Lag outcome ($t - 3$)	0.02 (-13.09, 13.12)	0.998
Lag outcome ($t - 4$)	-10.14 (-23.21, -2.91)	0.128
Lag outcome ($t - 5$)	-11.25 (-23.99, 1.49)	0.676

Abbreviation: CI, confidence interval.

* $p < 0.05$.*** $p < 0.001$.

significant interactions between catastrophizing scores and the first two lag terms ($t - 1$ and $t - 2$) in the switch/stay logistic model (Figure 2a), such that those with higher catastrophizing scores were more likely to stay after a win outcome, and switch after a loss outcome. There were no significant interactions with catastrophizing score in the reaction time model (Figure 2b). Our results indicate some interaction between history of outcomes and catastrophizing: we, therefore go on to report computational models that we designed to capture the relationship between outcome history and catastrophizing.

3.2 | Modeling results

Figure 3 demonstrates our fit assessment of six models to the data, according to the iBIC (Huys et al., 2012). Contrary to our prediction in H1, the best-fitting model was the standard Rescorla–Wagner Model (Rescorla, 1972), referred to henceforth as *1lr1t*. The *1lr1t* model equation is outlined below:

$$Q_{t+1,a} = Q_{t,a} + \text{learning rate} \cdot (\text{outcome}_t - Q_{t,a}), \quad (1)$$

where $Q_{t,a}$ is the learnt value of action a on trial t , *learning rate* is the extent to which learnt values are updated in response to reward

prediction errors (the part of the equation inside the brackets, representing the difference between the actual and expected outcomes), and outcome_t is the feedback from trial t . The free parameter per participant is *learning rate*.

We added an action model to the learning model which converts these learnt contingencies into choice probabilities. This action model includes an *inverse temperature* parameter, which governs choice noisiness and exploration, as shown below:

$$\text{Prob}_{t,a} = \frac{\exp(Q_{t,a} \cdot \text{inverse temperature})}{\sum(\exp(Q_{t,all} \cdot \text{inverse temperature}))}, \quad (2)$$

where *inverse temperature* is a parameter estimated for each participant and must have a positive value, and $Q_{t,all}$ signifies that the denominator term is a summation of exponent of the Q values for all stimuli on trial t (each multiplied by inverse temperature).

As stated in our preregistration (Harada-Laszlo et al., 2021), we additionally assessed the best-fitting model to include a *history* term. Our results showed the model in which *learning rate* is scaled by *history of negative outcome* (*lr scaling*), performed better than the learning model in which *sensitivity* is scaled by *history of negative outcome*, which aligns with our secondary prediction in H1. The *lr scaling* model equation is shown below:

$$Q_{t+1,a} = Q_{t,a} + (\text{learning rate} + (\text{scaling} \cdot \text{history}_t)) \cdot (\text{outcome}_t - Q_{t,a}), \quad (3)$$

where *scaling* represents a factor by which *learning rate* is adjusted and *history_t* is the proportion of trials that gave negative feedback within a window of the previous 10 trials (trial t inclusive). The free parameters are *learning rate* and *scaling*. We used the same action model as shown in Equation (2). Given the findings in the model agnostic analyses, we then created a family of models based on this one, with varying lengths of “window”: including 1–4 trials back (note that the last lag that was significant in the switch-stay analysis was 3, and the last lag in the RT analysis was 2). The best-fitting model had a “window” of 1 (Figure 3b), which in essence is somewhat similar in principle to a Vmax model—learning rates are adjusted when there is a change in the most rewarding outcome—whereas here, learning rates are adjusted just depending on whether there was negative feedback.

3.3 | Relationship between model parameters and catastrophizing symptoms

We performed correlation analyses between catastrophizing scores and the *scaling* parameter in the best model that included *history* (as in Hypothesis 2.1), as well as for all parameters in the winning model, as a predefined exploratory analysis (Harada-Laszlo et al., 2021). Catastrophizing had no relationship with the *scaling* parameter from the *lr scaling* model (Figure 4a), suggesting that catastrophizing does not relate to the extent to which individuals adjust their learning rate based on *history of negative outcome*. Additionally, our results demonstrated that neither the *learning rate*, nor the *inverse temperature*




FIGURE 2 Shows the results of mixed model analyses on switch/stay behavior and reaction time on trial t , with regressors including all significant lags from the previous models, and the interaction with catastrophizing scores. (a) mixed model of switch/stay behavior and (b) mixed model of reaction time.

parameter from the *1lr1t* model were significantly related to the propensity to catastrophise, indicating that catastrophizing is not related to the weighting of recent outcomes when updating learnt values (Figure 4b), nor exploration and stochasticity (Figure 4c).

3.4 | Multiple regression analyses

To examine whether other relevant variables affected our findings, we performed several planned multiple regression analyses to adjust for confounds and test specificity. As a reminder, the first category of these included relevant covariates (Model 1—working memory, fluid intelligence, age, and gender), and the second examined specificity to catastrophizing by including scores on other mental health questionnaires (Model 2).

3.5 | Scaling parameter (lr scaling model)

Catastrophizing was shown to have no effect on scaling in either regression model. Although the estimated coefficient value for

catastrophizing underwent a positive shift from Model 1 ($\beta = 0.001$, $p = 0.771$) to Model 2 ($\beta = 0.017$, $p = 0.095$) this did not reach significance at the predetermined $p = 0.05$ level (Figure 4d).

Notably, sex (where in the model we define *female* as the default state and effect is the result of being *male*) had a significantly negative effect on scaling in both Model 1 ($\beta = -0.466$, $p = 0.025$) and Model 2 ($\beta = -0.500$, $p = 0.017$), demonstrating that when the other predictors are controlled for, sex is negatively correlated with *scaling* by *history of negative outcome* (Figure 4d). This relationship was verified by further exploratory analysis (Supporting Information: Section 8.1)

3.6 | Learning rate (1lr1t model)

There was no significant relationship between catastrophizing and *learning rate* in either model. Similarly to the scaling parameter, the catastrophizing predictor underwent a positive shift from Model 1 ($\beta = 0.000$, $p = 0.960$), to Model 2 ($\beta = 0.003$, $p = 0.106$), which was ultimately insignificant (Figure 4e).

Fluid intelligence had a significant negative relationship with *learning rate* in both Model 1 ($\beta = -0.013$, $p = 0.015$) and Model 2

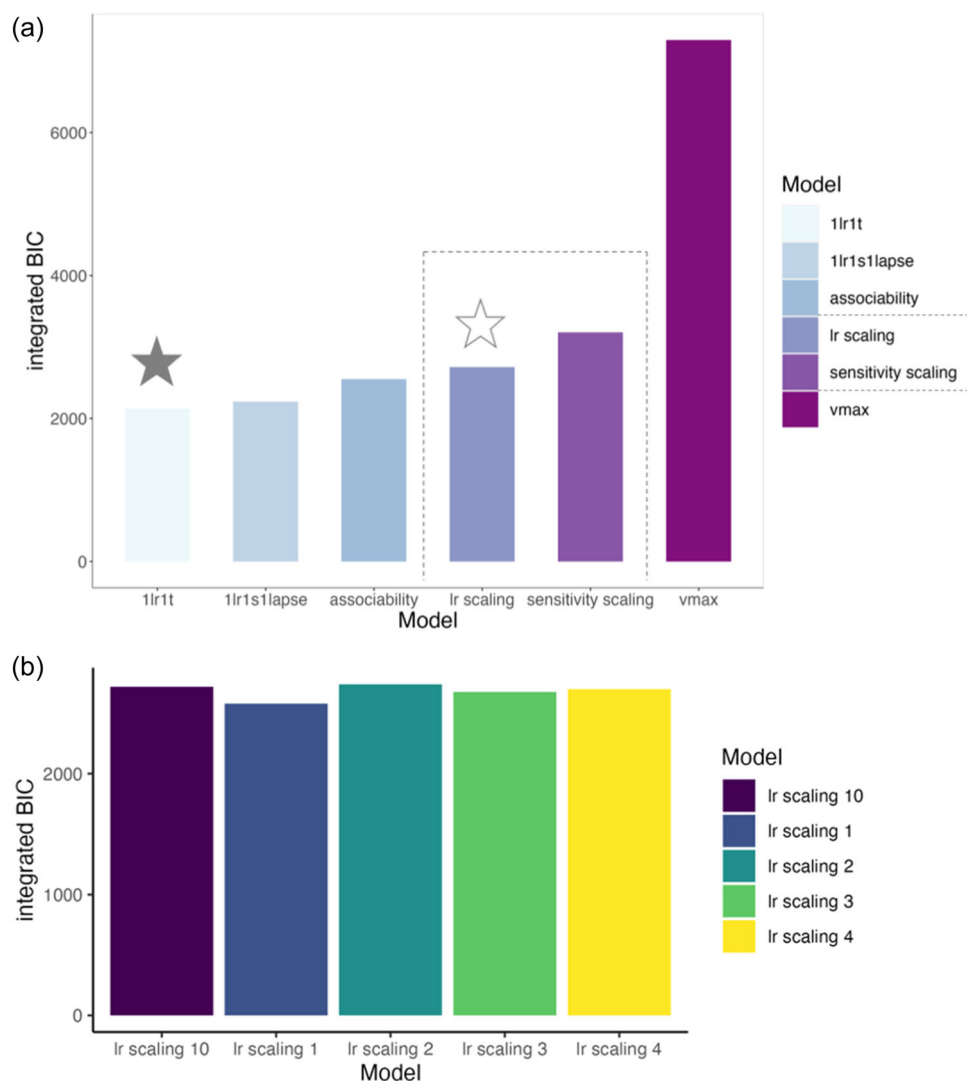


FIGURE 3 Illustrates the model fit results for our data as defined by the integrated Bayesian Information Criterion (iBIC). From best fit (lowest score) to worst fit (highest score). (a) initial six models, (b) family of models with scaled learning rates and different windows. 1lr1t = Standard Rescorla–Wagner, 1lr1s1lapse = Rescorla–Wagner model with sensitivity term, associability = Pearce–Hall model, lr scaling = learning rate scaled by history model (outcome window: 10), sensitivity scaling = sensitivity scaled by history model, vmax = Vmax model. In (a), the filled star denotes the winning model overall (1lr1t), while the outlined star marks the winning model out of those to include a history term (lr scaling).

($\beta = -0.013$, $p = 0.016$) indicating that greater fluid intelligence is associated with lower *learning rate*. We confirmed this effect in an exploratory analysis (Supporting Information: Section 8.2).

3.7 | Inverse temperature (1lr1t model)

Despite the figure displaying the catastrophizing coefficient having a weakly positive relationship with *inverse temperature* in Model 1 ($\beta = 0.004$, $p = 0.083$), this was not significant and the effect diminished in Model 2 ($\beta = 0.001$, $p = 0.901$) (Figure 4f).

Sex (where we define *female* as the default state and effect is the result of being *male*) had a strongly positive effect on *inverse temperature*, which was evident in both models [Model 1 ($\beta = 0.247$,

$p = 0.003$); Model 2 ($\beta = 0.246$, $p = 0.004$)], while age was negatively related to *inverse temperature* in both models [Model 1 ($\beta = -0.009$, $p = 0.011$); Model 2 ($\beta = -0.009$, $p = 0.020$)]. These relationships were supported by exploratory analysis (Supporting Information: Section 8.3), indicating that in our data, *inverse temperature* is higher in men than in women and decreases with age.

4 | DISCUSSION

In this study, we explored the relationship between reinforcement learning parameters and self-report catastrophizing scores to understand how learning processes may be transformed in catastrophizing, with a particular interest in the *scaling by history of negative outcome* parameter,

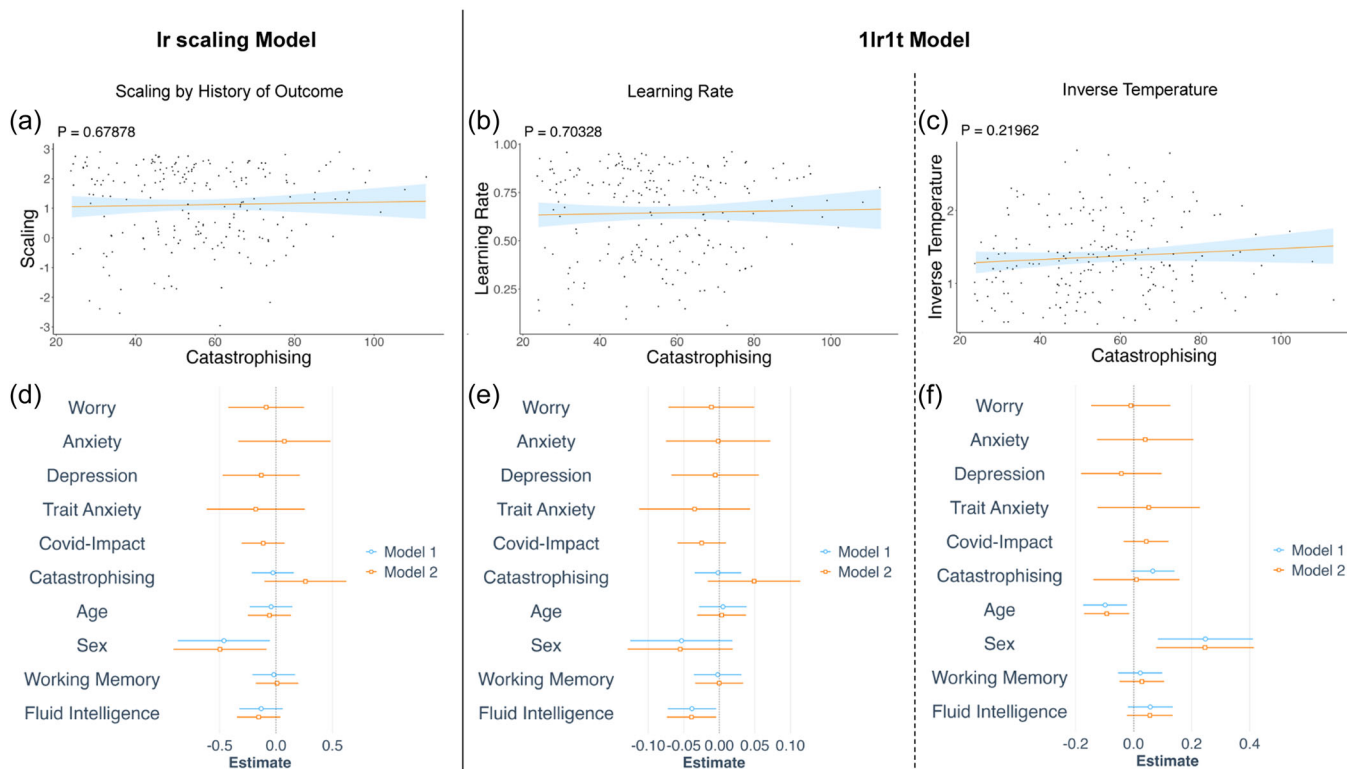


FIGURE 4 (a–c) Illustrate the plotted correlations between self-report catastrophizing scores and specific reinforcement learning parameters. None of these relationships were found to be statistically significant; (a) shows the relationship between catastrophizing and scaling ($r = 0.003$, $p = 0.679$), (b) with learning rate ($r = 0.026$, $p = 0.703$), and (c) with inverse temperature ($r = 0.086$, $p = 0.220$). (d–f) Show multiple regression analyses for each parameter, with estimated coefficients of all Model 1 (blue) and Model 2 (orange) predictors displayed on the same plot to convey the effects of adding questionnaires as covariates. (d) Shows multiple regression analysis for the scaling parameter. Catastrophizing had no effect on the scaling parameter in Model 1 ($\beta = -0.001$, $p = 0.771$) or Model 2 ($\beta = 0.017$, $p = 0.095$). There was a significant negative effect of self-reported sex on the scaling parameter in both Model 1 ($\beta = -0.466$, $p = 0.025$) and Model 2 ($\beta = -0.500$, $p = 0.017$). (E) Shows multiple regression analysis for the learning rate parameter of the 1lr1t model. Catastrophizing had no effect in Model 1 ($\beta = 0.000$, $p = 0.960$) or Model 2 ($\beta = 0.003$, $p = 0.106$). Fluid intelligence had a significant negative effect on learning rate parameter in Model 1 ($\beta = -0.013$, $p = 0.015$) and Model 2 ($\beta = -0.013$, $p = 0.016$). (F) Shows multiple regression analysis for the inverse temperature parameter of the 1lr1t model. Catastrophizing had no significant effect on Model 1 ($\beta = 0.004$, $p = 0.083$) or Model 2 ($\beta = 0.001$, $p = 0.901$). Sex (male) had a strong positive effect on inverse temperature in Model 1 ($\beta = 0.247$, $p = 0.003$) and Model 2 ($\beta = 0.246$, $p = 0.004$), while age had a strong negative effect in Model 1 ($\beta = -0.009$, $p = 0.011$) and Model 2 ($\beta = -0.009$, $p = 0.020$).

which we considered to be a promising construct to relate to catastrophizing (H2.1). First, we performed model-agnostic analyses, which showed that history of outcomes does impact future behavior (positive outcomes promote “stay” behavior, and neutral outcomes promote “switch” behavior for three trials into the future, and positive outcomes promote an initial slowing of response, then subsequent speeding), and that (at least in our “accuracy” analyses), this may be enhanced by catastrophizing. We then conducted model fitting on several standard learning models that are prevalent in computational psychiatry, in addition to two novel models which were designed to explicitly parameterize altered learning when a series of negative outcomes are received. Our model comparison showed that the 1lr1t model, which did not include a *scaling* term, was the best fit for our data overall, and the *Ir scaling* model (with a “history” window of only one trial) was the best-fitting model to include a *history* variable.

We did not find a relationship between catastrophizing and any of the parameters we investigated. However, we did find significant

relationships between parameters and several of our covariates. In the *Ir scaling* model, sex (where we define *female* as the default state and effect is the result of being *male*) was negatively related to the *scaling* parameter. In the 1lr1t model, fluid intelligence was negatively correlated with the *learning rate* parameter, while the *inverse temperature* parameter was negatively correlated to age, and positively related to sex.

4.1 | Standard reinforcement learning models performed better than our novel models with a scaling parameter

A notable finding from this study is that the Standard Rescorla–Wagner model (Rescorla, 1972) (1lr1t model) provided a better fit for the data than our models which included a scaling term related to the history of outcome. It can, therefore, be inferred that in

this task, the participant learning rate was unaffected (or only minimally affected) by the proportion of negative outcomes received. It is important to note that prediction errors will increase with repeated incorrect choices, and thus scaling may only have a minor additional effect above that of prediction errors. Our study additionally assessed the fit of the Vmax model (Supporting Information: Section 6), whereby learning rate increases if the most valuable option has changed from the previous trial. Given the performance of this model was inferior even to the novel models, it suggests that proportional adjustments in learning rate based on poor previous outcome may not be a feature of human learning and decision making, at least on this task.

4.2 | Catastrophizing has no relationship with reinforcement learning parameters

Counter to our predictions, we were unable to detect a relationship between self-report catastrophizing scores and any one of the reinforcement learning parameters we explored in this study, demonstrating that, within our sample, catastrophizing had no effect on *learning rate*, *inverse temperature*, or the extent to which individuals may scale *learning rate* by *history of negative outcome*. Notably, inclusion of mental health covariates in a multiple regression analysis led to a trend in the catastrophizing coefficient toward the predicted positive direction for both the *scaling* parameter of the *lr scaling* model and the *learning rate* parameter of the *1lr1t* model. Though neither reached significance at the predetermined level, the common pattern across models may imply that the study was underpowered to capture this effect, or that there are confounding variables masking an effect of catastrophizing. Future replication with a larger sample size would be required to confirm this.

Alternative explanations of our failure to detect a relationship include the possibility that a more complex relationship exists between catastrophizing and reinforcement learning parameters, whereby individual differences in the effects that catastrophizing elicit on behavior evoke contradictory or inconsistent behavioral responses amongst participants. In this case, we may not have detected an effect, as correlations and regressions assume linear relationships without individual differences in the direction of the relationship. Latent class analysis or correlation clustering could be utilized to determine if there is evidence for individual differences in response. Moreover, it is feasible that there are valence-dependent effects on learning in catastrophizing, which we did not assess, due to our models including single learning rates, thus not allowing the separation of learning from reward and learning from punishment. Such effects have been shown in mood and anxiety disorders, which feature considerable comorbidity with catastrophizing (Aylward et al., 2019).

4.3 | Significant covariate effects on reinforcement learning parameters

We observed significant (unhypothesised) relationships between our selected reinforcement learning parameters and several of our

covariates, which are individually discussed in the Supporting Information as secondary findings (Supporting Information S1: Section 9). Beyond the individual implications of each of these findings, the inherent presence of significant effects suggests that the selected reinforcement learning parameters are likely to be meaningfully capturing a cognitive construct, but that these constructs were unrelated to the mental health measures collected in this study.

4.4 | Win-stay behavior and catastrophizing

Our model-agnostic analyses indicate that there may be some relationship between win-stay behavior and catastrophizing, which was not captured in our computational modeling analyses. Future research could develop clear hypotheses and models that might allow this behavior to be captured and mechanistically characterized—perhaps this reflects a greater tendency to “perseverate” in those who catastrophize, which may be captured by a perseveration parameter in a model. Alternatively, perseveration after win outcomes may reflect a form of “safety behavior” or avoidance, which could perhaps be captured by a go/no-go paradigm modeled using an avoidance parameter (Guitart-Masip et al., 2012; Mkrtchian et al., 2017).

5 | LIMITATIONS

Notably, in this bandit task the possible outcomes were rewards (+1 points) and neutral outcomes (0 points). Despite our attempt to ensure that neutral outcomes were perceived as negative (participants received a message saying that the restaurant choice was “disgusting,” and fewer points won translated to a lower financial bonus), it may be more relevant to catastrophizing and other mental health constructs to use explicit punishments (loss of points, rather than absence of gain of points or financial bonus) in future research examining history of outcomes.

Moreover, it is possible that catastrophizing is only present or primarily present during high arousal or anxiety (Clark et al., 1988), or that it requires activation by a trigger specific to the individual (Beck, 1976). As we did not manipulate arousal or anxiety, and the stakes of our gamified abstract task were both non-threatening and impersonal, it is unlikely that our task evoked elevated anxiety in our participants, thus potentially limiting their catastrophizing behavior.

6 | CONCLUSIONS

We were unable to provide any conclusive evidence as to whether catastrophizing relates to reinforcement learning parameters. Nevertheless, there were some ambiguous trends in the effect of catastrophizing on reinforcement learning parameters in a direction that aligned with our predictions. However, our model fit findings suggest that learning rate may not be proportionally affected by the

history of outcomes beyond the effect of prediction error, and that standard reinforcement learning frameworks may be superior in capturing participant learning and decision making on this task, over our novel model which included a scaling term.

ACKNOWLEDGMENTS

O. J. R.'s MRC senior fellowship is partially in collaboration with Cambridge Cognition (who plan to provide in-kind contribution) and he is running an investigator-initiated trial with medication donated by Lundbeck (escitalopram and placebo, no financial contribution). He also holds an MRC-Proximity to discovery award with Roche (who provide in-kind contributions and have sponsored travel for A. C. P.) regarding work on heart rate variability and anxiety. He has also completed consultancy work on affective bias modification for Peak and online CBT for IESO digital health. A. C. P. has received funding from the Wellcome Trust (Grant ref: 226694/Z/22/Z). M. H. L. received funding from UCL Medical Sciences Department to undertake a research project as part of the iBSc degree requirements. A. T. currently works for Haleon, as a data scientist. The authors declare no other conflicts of interest. The MRC (who provided funding) and Haleon had no role in the study design, collection, analysis, or interpretation of data.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available on the OSF at [10.17605/OSF.IO/3Y6UJ](https://doi.org/10.17605/OSF.IO/3Y6UJ). This study was preregistered: [10.17605/OSF.IO/VSN8B](https://doi.org/10.17605/OSF.IO/VSN8B).

ORCID

Alexandra C. Pike  <https://orcid.org/0000-0003-1972-5530>

REFERENCES

- Austin, D. W., & Richards, J. C. (2001). The catastrophic misinterpretation model of panic disorder. *Behaviour Research and Therapy*, 39(11), 1277–1291.
- Averbeck, B. B. (2015). Theory of choice in bandit, information sampling and foraging tasks. *PLoS Computational Biology*, 11(3), e1004164.
- Aylward, J., Valton, V., Ahn, W. Y., Bond, R. L., Dayan, P., Roiser, J. P., & Robinson, O. J. (2019). Altered learning under uncertainty in unmedicated mood and anxiety disorders. *Nature Human Behaviour*, 3(10), 1116–1123. <https://doi.org/10.1038/s41562-019-0628-0>
- Beck, A. T. (1976). *Cognitive therapy and the emotional disorders*. International Universities Press.
- Bowers, W. A., Evans, K., & Andersen, A. E. (1997). Inpatient treatment of eating disorders: A cognitive therapy milieu. *Cognitive and Behavioral Practice*, 4(2), 291–323.
- Chierchia, G., Fuhrmann, D., Knoll, L. J., Pi-Sunyer, B. P., Sakhardande, A. L., & Blakemore, S.-J. (2019). The matrix reasoning item bank (MaRs-IB): Novel, open-access abstract reasoning items for adolescents and adults. *Royal Society Open Science*, 6(10), 190232.
- Clark, D., Salkovskis, P., Gelder, M., Koehler, C., Martin, M., Anastasiades, P., Hackmann, A., Middleton, H., & Jeavons, A. (1988). Tests of a cognitive theory of panic. In I. Hand & H. U. Wittchen (Eds.), *Panic and phobias* (Vol. 2, pp. 149–158). Springer.
- Clark, D. A., & Beck, A. T. (2011). *Cognitive therapy of anxiety disorders: Science and practice*. Guilford Press.
- Collins, A. G. E., Brown, J. K., Gold, J. M., Waltz, J. A., & Frank, M. J. (2014). Working memory contributions to reinforcement learning impairments in schizophrenia. *The Journal of Neuroscience*, 34(41), 13747–13756.
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876–879.
- Dayan, P. (2009). Dopamine, reinforcement learning, and addiction. *Pharmacopsychiatry*, 42(S 01), S56–S65.
- Dunmore, E., Clark, D. M., & Ehlers, A. (1999). Cognitive factors involved in the onset and maintenance of posttraumatic stress disorder (PTSD) after physical or sexual assault. *Behaviour Research and Therapy*, 37(9), 809–829. [https://doi.org/10.1016/s0005-7967\(98\)00181-8](https://doi.org/10.1016/s0005-7967(98)00181-8)
- Ellis, A. (1962). *Reason and emotion in psychotherapy*. Lyle Stuart.
- Gellatly, R., & Beck, A. T. (2016). Catastrophic thinking: A transdiagnostic process across psychiatric disorders. *Cognitive Therapy and Research*, 40(4), 441–452.
- Guitart-Masip, M., Huys, Q. J. M., Fuentemilla, L., Dayan, P., Duzel, E., & Dolan, R. J. (2012). Go and no-go learning in reward and punishment: Interactions between affect and effect. *NeuroImage*, 62(1), 154–166. <https://doi.org/10.1016/j.neuroimage.2012.04.024>
- Harada-Laszlo, M., Pike, A. C., Talwar, A., & Robinson, O. J. (2021). *Exploring how learning is implicated in catastrophizing by modelling a computerised task*. OSF Preregistration.
- Hoffman, M. D., & Gelman, A. (2014). The no-u-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15, 1593–1623.
- Huang, H., Thompson, W., & Paulus, M. P. (2017). Computational dysfunctions in anxiety: Failure to differentiate signal from noise. *Biological Psychiatry*, 82(6), 440–446. <https://doi.org/10.1016/j.biopsych.2017.07.007>
- Huys, Q. J., Pizzagalli, D. A., Bogdan, R., & Dayan, P. (2013). Mapping anhedonia onto reinforcement learning: A behavioural meta-analysis. *Biology of Mood & Anxiety Disorders*, 3(1), 12. <https://doi.org/10.1186/2045-5380-3-12>
- Huys, Q. J. M., Eshel, N., O'Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: How the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, 8(3), e1002410.
- Jenness, J. L., Jager-Hyman, S., Heleniak, C., Beck, A. T., Sheridan, M. A., & McLaughlin, K. A. (2016). Catastrophizing, rumination, and reappraisal prospectively predict adolescent PTSD symptom onset following a terrorist attack. *Depression and Anxiety*, 33(11), 1039–1047.
- Kroenke, K., Spitzer, R. L., & Williams, J. B. W. (2001). The PHQ-9: Validity of a brief depression severity measure. *Journal of General Internal Medicine*, 16(9), 606–613. <https://doi.org/10.1046/j.1525-1497.2001.016009606.x>
- Kroenke, K., Strine, T. W., Spitzer, R. L., Williams, J. B. W., Berry, J. T., & Mokdad, A. H. (2009). The PHQ-8 as a measure of current depression in the general population. *Journal of Affective Disorders*, 114(1–3), 163–173. <https://doi.org/10.1016/j.jad.2008.06.026>
- Kumar, P., Waiter, G., Ahearn, T., Milders, M., Reid, I., & Steele, J. D. (2008). Abnormal temporal difference reward-learning signals in major depression. *Brain*, 131(8), 2084–2093.
- Lawson, R. P., Bisby, J., Nord, C. L., Burgess, N., & Rees, G. (2021). The computational, pharmacological, and physiological determinants of sensory learning under uncertainty. *Current Biology*, 31(1), 163–172.e4. <https://doi.org/10.1016/j.cub.2020.10.043>
- McLaughlin, K. A., Busso, D. S., Duys, A., Green, J. G., Alves, S., Way, M., & Sheridan, M. A. (2014). Amygdala response to negative stimuli predicts PTSD symptom onset following a terrorist attack. *Depression and Anxiety*, 31(10), 834–842.

- Meyer, T. J., Miller, M. L., Metzger, R. L., & Borkovec, T. D. (1990). Development and validation of the Penn State Worry Questionnaire. *Behaviour Research and Therapy*, 28(6), 487–495. [https://doi.org/10.1016/0005-7967\(90\)90135-6](https://doi.org/10.1016/0005-7967(90)90135-6)
- Mkrtchian, A., Aylward, J., Dayan, P., Roiser, J. P., & Robinson, O. J. (2017). Modeling avoidance in mood and anxiety disorders using reinforcement learning. *Biological Psychiatry*, 82(7), 532–539. <https://doi.org/10.1016/j.biopsych.2017.01.017>
- Paré, C., Thibault, P., Côté, P., Pimentel, S. D., Shen, S., Yakobov, E., & Sullivan, M. J. L. (2019). The relationship between level of catastrophizing and mental health comorbidity in individuals with whiplash injuries. *The Clinical Journal of Pain*, 35(11), 880–886.
- Park, C., Pan, Z., Brietzke, E., Subramaniapillai, M., Rosenblat, J. D., Zuckerman, H., Lee, Y., Fus, D., & McIntyre, R. S. (2018). Predicting antidepressant response using early changes in cognition: A systematic review. *Behavioural Brain Research*, 353, 154–160. <https://doi.org/10.1016/j.bbr.2018.07.011>
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87(6), 532–552.
- Pike, A. C., Alves Anet, Á., Peleg, N., & Robinson, O. J. (2023). Catastrophizing and risk-taking. *Computational Psychiatry*, 7, 1. <https://doi.org/10.5334/cpsy.91>
- 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.
- Pike, A. C., Serfaty, J. R., & Robinson, O. J. (2021). The development and psychometric properties of a self-report Catastrophizing Questionnaire. *Royal Society Open Science*, 8(1), 201362.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). Appleton Century Crofts.
- Schad, D. J., Jünger, E., Sebold, M., Garbusow, M., Bernhardt, N., Javadi, A.-H., Zimmermann, U. S., Smolka, M. N., Heinz, A., Rapp, M. A., & Huys, Q. J. M. (2014). Processing speed enhances model-based over model-free reinforcement learning in the presence of high working memory functioning. *Frontiers in Psychology*, 5, 1450.
- Smith, E. M., Reynolds, S., Orchard, F., Whalley, H. C., & Chan, S. W. (2018). Cognitive biases predict symptoms of depression, anxiety and wellbeing above and beyond neuroticism in adolescence. *Journal of Affective Disorders*, 241, 446–453. <https://doi.org/10.1016/j.jad.2018.08.051>
- Spielberger, C. D. (1983). *State-trait anxiety inventory for adults*. Consulting Psychologists Press.
- Spitzer, R. L., Kroenke, K., Williams, J. B. W., & Löwe, B. (2006). A brief measure for assessing generalized anxiety disorder: The GAD-7. *Archives of Internal Medicine*, 166, 1092–1097.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning* (Vol. 135). MIT Press.
- Wechsler, D. (1955). *Wechsler adult intelligence scale*. Psychological Corp.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Harada-Laszlo, M., Talwar, A., Robinson, O. J., & Pike, A. C. (2023). A series of unfortunate events: Do those who catastrophize learn more after negative outcomes? *Mental Health Science*, e49. <https://doi.org/10.1002/mhs2.49>