

Computational Modelling of Childhood Trauma (Complex PTSD)

Jeffrey Scholes



Master of Science
Data Science
School of Informatics
University of Edinburgh
2023

Abstract

Post-traumatic stress disorder (PTSD) is a psychiatric disorder that can develop after experiencing traumatic, usually life-threatening events. The disorder is characterised through various symptom clusters, each of which can cause extreme distress and result in significant impairment within many aspects of an individuals life. Complex PTSD (cPTSD) is a recently proposed distinction of the disorder, resulting from prolonged, ongoing trauma that can cause additional symptoms, and typically precipitates lasting damage to an individuals perception of themselves and the world. Current knowledge of the underlying mechanisms of this disorder remains limited, as does knowledge of clear distinctions between PTSD types. Through methods of computational psychiatry, we model the fear learning aspect of the disorder, hoping to improve this understanding. Several reinforcement learning models of PTSD are summarised. A recent TD-Momentum model is then re-implemented and evaluated in terms of describing cPTSD and implications to treatments. Several extensions are then proposed that can add value to the model, these include concepts of associability, outcome-sensitivity, and valence partitioning which are evaluated in terms of their implications.

Research Ethics Approval

This project was planned in accordance with the Informatics Research Ethics policy. It did not involve any aspects that required approval from the Informatics Research Ethics committee.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Jeffrey Scholes)

Acknowledgements

First of all, I would like to extend immense thanks to my supervisor, Dr. Peggy Seriès, for the continued support, valuable feedback, and for introducing me to computational psychiatry. I am extremely grateful to have had the opportunity to make my own contribution to this very meaningful field under your excellent guidance.

I am also very thankful to Filippo Ferrari for the extended support and expert insight throughout this project.

I would like to take this opportunity to thank Dr. Iain Murray, whose kind words at a time when I was struggling this year meant a great deal to me.

Finally, I am extremely grateful for the love and support of all my family and friends, particularly my mum and dad.

Table of Contents

1	Introduction	1
2	Background	3
2.1	Literature Review - Reinforcement Learning Models	4
2.1.1	Kaye et al. (2023) - TD-Momentum Threat Prediction Model	5
2.1.2	Homan et al. (2019) & Brown et al. (2018) - Associability Models of Fear Learning and Hyperarousal	5
2.1.3	Ross et al. (2018) - General PTSD Learning Deficits	7
2.1.4	Yanamori et al. (2023) - Anxiety Related Approach-Avoidance	8
2.1.5	Eldar et al. (2016) - Momentum Model of Mood in Bipolar Disorder	9
3	Re-implementation & Evaluation of Kaye et al. (2023)	11
3.1	Bayesian Baseline Model	12
3.1.1	Bayesian Attack Rate Estimation Model	12
3.1.2	Auto-regressive Time Series	13
3.1.3	Bayesian Model - Re-implementation	14
3.1.4	Bayesian Model - Additional Investigations	15
3.2	Stress-Enhanced Fear-Learning Model	17
3.3	Reinforcement Learning (RL) Models	18
3.3.1	Parameter Recovery	18
3.3.2	TD Model	19
3.3.3	TD Model - Re-implementation	20
3.3.4	TD-Momentum Model	21
3.3.5	TD-Momentum Model - Re-implementation	21
3.3.6	TD-Momentum Model - Additional Investigations	23

4	Extending the TD-Momentum Model of Kaye et al. (2023)	29
4.1	Incorporating Associability in TD-Momentum Model	29
4.1.1	Associability TD-Momentum Model Form	30
4.1.2	Implications of Associability TD-Momentum	30
4.2	Incorporating Outcome-Sensitivity in TD-Momentum Model	32
4.2.1	Outcome-Sensitive TD-Momentum Model Form	32
4.2.2	Implications of Outcome-Sensitive TD-Momentum	33
4.3	Incorporating Valence-Partitioning in TD-Momentum Model	34
4.3.1	Valence-Partitioned TD-Momentum Model Form	35
4.3.2	Implications of Valence-Partitioned TD-Momentum	36
4.4	TD-Momentum Extensions Summary	38
5	Conclusion	39
	Bibliography	41
A	Additional Plots	45
B	Bayesian Model Python Code	52
B.1	Bayesian Model - Attack & Death Simulations	52
B.2	Bayesian Model - MCMC Sampler	53
B.3	Bayesian Model - Posterior Plots	56
B.4	Bayesian Model - Variance in Threat Over Time	58
B.5	Bayesian Model - AR Time Series	59
B.6	Bayesian Model - Posterior Comparisons for Various c Coefficients in AR Time Series	60
B.7	Bayesian Model - ELS Example	62
B.8	Bayesian Model - Clustered ELS Attacks	64
C	TD Model Python Code	67
C.1	Preprocessing SEFL Data	67
C.2	TD Model - Smoothed Freezing & Fitting SEFL Data	71
C.3	TD Model - Computing BIC Scores	84
C.4	TD Model - Multivariate Sampled Parameter Recovery	86
C.5	TD Model - List Sampled Parameter Recovery	91

D	TD-Momentum Model Code	93
D.1	TD-Momentum Model - Example of SEFL Model Fit & Related Freezing (Fig 3.7)	93
D.2	TD-Momentum Model - Fitting SEFL Data	100
D.3	TD-Momentum Model - Computing BIC Scores	124
D.4	TD-Momentum Model - Multivariate Sampled Parameter Recovery	125
D.5	TD-Momentum Model - List Sampled Parameter Recovery	132
D.6	BIC Model Comparison	134
E	Extension Simulations Code	136
E.1	Associability TD-Momentum Model	136
E.2	Risk-Sensitive TD-Momentum Model	141
E.3	Valence-Partitioned TD-Momentum Model	145

Chapter 1

Introduction

Post-traumatic stress disorder (PTSD) is a psychiatric disorder which affects all ages and can develop after an individual experiences or witnesses a traumatic event. Although not all who suffer from PTSD have experienced a dangerous life event, the risk of developing PTSD drastically increases after exposure to serious trauma. The context of such traumatic events can vary greatly, but will likely include feelings of horror, extreme fear or helplessness, e.g. being injured or witnessing another person being injured or killed. The disorder is thought to affect around 6–8% of the general population, with this figure increasing significantly up to 25% if individuals have experienced severe psychological trauma (combat veterans, war refugees) [22].

The Diagnostic and Statistical Manual of Mental Disorders (DSM-5) [4] classifies PTSD as a Trauma- and Stressor-Related Disorder, rather than an anxiety disorder and organises symptoms into clusters, where the duration of disturbances must be more than one month and cause significant distress and/or impairment to an individual's social or occupational life (which cannot be attributed to effects of substance abuse or another medical condition). Symptom clusters include: **intrusion symptoms**, **avoidance symptoms**, **negative alterations in cognition/mood**, and **alterations in arousal and reactivity**. However, the DSM-5 does not separate type-1 and type-2 PTSD (also known as complex PTSD, or cPTSD), so no clinical diagnosis of cPTSD currently exists. cPTSD refers to PTSD caused by repeated, long-lasting exposure to traumatic experiences (at any age) that may not be life threatening themselves, but are still extremely traumatic and often come with an increased sense of helplessness and powerlessness (e.g. domestic abuse, sexual abuse).

In contrast to the DSM-5, the International Classification of Diseases (ICD-11) [29] does provide a distinction for PTSD types. The ICD-11 lists PTSD symptoms

to be: **re-experiencing symptoms** (intrusive, distressing memories, dissociative reactions e.g. flashbacks, intense distress at exposure to cues reminiscent of the trauma, re-experiencing trauma in the here and now), **avoidance symptoms** (deliberate avoidance of memories and external reminders related to trauma, e.g. people and places), **alterations in arousal and reactivity** (emotional outbursts, persistent perceptions of heightened fear, hypervigilance, hyperarousal). Additional clinical features are noted, including dissociative symptoms, suicidal ideation, substance abuse, social withdrawal, and anxiety. Emotional responses include sadness, shame, humiliation, and guilt (also survivor guilt). The criteria for cPTSD includes all of the above, and additionally: **severe problems with affect regulation** (increased emotional reactivity, self-destructive behaviour, emotional numbing, dissociation), **persistent diminished feelings of oneself** (worthlessness, guilt from not having escaped the trauma/preventing it for others), **persistent difficulties in maintaining relationships** (avoiding relationships/social situations, a sustained view of the world being dangerous/that people cannot be trusted).

There is much debate as to whether these PTSD types should be distinctly separated. Some, like Cloitre et al. [8], believe that a clear distinction is necessary for progress in refining treatments and diagnoses. However, there are those, such as Achterhof et al. [2], who claim that statistical tests show that patient groups are not well separated, showing a lack of clear evidence to support such a distinction. Here, we take the stance of Cloitre et al., that a distinction between PTSD and cPTSD will benefit more tailored and effective treatment planning. Due to this distinction being a recent inclusion in the ICD-11, there is much ambiguity in the literature with regards to which PTSD type is being discussed, making it difficult to clearly identify if the literature is more relevant to cases of prolonged, ongoing trauma or sudden, more classical PTSD trauma.

Ultimately, this dissertation aims to provide insight on cPTSD processes and symptoms, and to provide avenues for further research that could help in progressing treatment. Our focus being on why treatments may not always be effective (treatment failure/dropout), and how cPTSD mechanisms may differ to more classical PTSD. We begin by reviewing typical approaches of computational PTSD models, we then review the literature of relevant models, discussing what they can explain and their shortcomings. We then re-implement the main models of Kaye et al. [16], reproducing results and noting the extent to which they can explain mechanisms of cPTSD. We also perform our own additional investigations on how model behaviour reacts to early life stress (ELS). Finally, we provide details on potential model extensions that can add value to explaining mechanisms of the disorder.

Chapter 2

Background

Many key questions remain in order for PTSD to be better understood. The underlying mechanisms of the disorder need to be clarified for treatments of each PTSD type to be improved, and for clinicians to be able to accurately predict which individuals are more susceptible to developing the disorder after experiencing trauma. Some important research questions that remain include [5]:

- Why (at a neurobiological level) do some individuals fully recover from traumatic experiences and others do not?
- Can we predict how different people will react to trauma and/or treatment?
- How does severity and duration of trauma affect likelihood and severity of PTSD?

There have been various different approaches taken in modelling PTSD mechanisms and framing these important questions. Much of the literature takes on one or more of these approaches: predictive coding, Bayesian inference, and reinforcement learning.

The predictive coding framework views the brain as a predicting machine that minimises prediction errors to correctly update predictions based on sensory input. PTSD is proposed to cause a breakdown in this framework that is reflected in symptom clusters. Aitchison and Lengyel note that while predictive coding offers a framework for implementing prediction error based learning updates, Bayesian inference provides a possible calculus for computing such predictions [3]. In Bayesian inference, prior beliefs about the environment/sensory input are tested and updated when new evidence is received, where the posterior updates follow Bayes theorem. The prior belief is combined with the likelihood of observed data given the prior belief, and is normalised by the overall probability of observing the data, giving a valid posterior. This process repeats using the latest posterior as the new prior, generating an endless cycle of belief

updates about the environment. Intense trauma can change an individual's internal model such that powerful, maladaptive priors reflect PTSD symptom clusters.

Reinforcement learning (RL) models generate agents that learn a set task in real time through trial and error, along with delayed rewards. They are typically applied to model learning or decision making, e.g. avoidance or fear learning and are well suited to be fit to behavioural tasks, with their adaptability allowing us to identify differences in the neural functions of learning processes in individuals with PTSD. These models are typically combined with neuroimaging data from the areas of the brain associated with learning and decision-making (e.g. the amygdala), allowing for a deeper understanding of the mechanistic processes of learning related to the disorder [7]. However, current RL approaches are usually hindered by basic assumptions of outcomes in the related behavioural tasks, which may not account for the broad spectrum of outcomes in reality.

Predictive coding approaches to PTSD are typically theoretical and not applied to data due to the complexity of retrieving such information (e.g. how do we extract accurate information about one's internal belief system?). However, they do provide a thorough base for theory on the underlying mechanisms of PTSD and can be applied to explain a variety of symptom clusters [30]. The following literature review focuses on RL models of PTSD that relate more directly, and can be compared to, the main RL model explored in this dissertation.

2.1 Literature Review - Reinforcement Learning Models

Radell et al. suggest that a fully comprehensive model of PTSD should provide novel predictions and potential explanations for the underlying mechanisms of *all* symptom clusters of the disorder [19]. However, the complexity of the disorder, along with the current lack of understanding the underlying mechanisms, make this unifying model difficult to define. Most computational models of PTSD instead focus on adaptations within one of the symptom clusters: fear learning, hyperarousal, avoidance, cognition & mood, and intrusive memories. With RL models commonly applied to behavioural data, symptom clusters where it is easier to extract such data are favoured, namely: fear learning, hyperarousal, and avoidance. Models related to disorders sharing symptomatology with PTSD are included from which we can draw parallels to PTSD mechanisms. The lack of distinction between PTSD types within the literature makes it difficult to specify which models relate to cPTSD. Efforts have been made to include relevant literature with the potential to describe mechanisms associated with cPTSD.

2.1.1 Kaye et al. (2023) - TD-Momentum Threat Prediction Model

Kaye et al. [16] propose a fear learning model of estimating threat which incorporates associative and non-associative responses to threat across different contexts. They implement a Temporal Difference learning model (TD) with the addition of a *momentum* term, m_t . Threat in context c at time t is computed as follows.

$$T_{c,t} = T_{c,t-1} + \alpha(u_t - \gamma_1 T_{c,t-1}) + f m_t \quad (2.1)$$

$$m_t = m_{t-1} + \gamma_2 \sum_{c=\{A,B,\dots\}} \alpha(u_t - T_{c,t-1}) \quad (2.2)$$

γ_1 is the decay rate of associative prediction errors, α is the learning rate. The momentum at time t , m_t , is computed from the sum of all decayed prediction errors across all contexts. Scaling constant f and momentum decay rate γ_2 control how much momentum influences learning.

Momentum was first proposed by Eldar et al. [10] as a way of representing mood in Bipolar Disorder (section 2.1.5). Eldar et al. proposed that mood corresponds to the overall momentum of recent outcomes, where its biasing influence on the perception of outcomes accounts for environmental dependencies, therefore “correcting” learning. Applied here by Kaye et al., momentum (Eq. 2.2) allows for threat prediction errors in any given context to affect predictions in all contexts. In other words, this term corresponds to the “mood” of an agent, given recent experiences.

Comparing to a basic TD model and fitting to real data from a contextual rodent study (discussed in chapter 3), the authors showed how this TD-Momentum model represented fear learning better than the simpler model. Kaye et al. also note the addition of this momentum term may provide potential reasons for PTSD treatment failure (exposure therapy). However, the model is only evaluated generally by the authors. In chapter 3 we re-implement the models and evaluate them with a focus on how early life stress can effect threat perception. We also explore any links to cPTSD.

2.1.2 Homan et al. (2019) & Brown et al. (2018) - Associability Models of Fear Learning and Hyperarousal

Homan et al. create a fear learning hybrid model of Rescorla-Wagner/Pearce-Hall models [14] used to explain conditioned threat responses for two groups of combat veterans (PTSD and healthy controls). The task for both groups was to learn pairings of pictures of faces and electric shocks. The authors identified a correlation between

the magnitude of prediction errors and the severity of symptoms in PTSD individuals, where highly symptomatic individuals were more influenced by larger prediction errors.

The associability term gives a value to the attention that each specific cue receives, depending on its historic accuracy of predicting the outcome. Associability effectively turns the learning rate from a constant into a dynamic parameter. Unreliable cues receive larger associability as time moves on, as they are more likely to be unreliable in the future, and are thus updated preferentially as new information comes in.

Brown et al. [7] propose a similar hybrid Rescorla-Wagner/Pearce-Hall model, focusing on hyperarousal symptoms and investigating potential reasons for disproportionate reactions to unexpected stimuli. The authors fit results from a loss learning task performed by a group of combat veterans who had to choose between two stimuli with monetary income and learn the “better” option over time. Brown et al. showed that veterans with PTSD have an increased learning response to surprising events, i.e. an increased learning rate for unexpected events that inherently cause larger prediction errors. PTSD individuals typically received increased associability weights for unexpected cues during loss learning, meaning that loss learning in these individuals was more heavily influenced by unexpected outcomes related to these cues. Similar to Homan et al., associability represents the attention that each cue is given by the participant. Thus, loss learning was more heavily influenced by attention given to surprising outcomes, with PTSD individuals more likely to allocate additional attention to unexpected outcomes, which could provide reason for exaggerated responses to unexpected stimuli (hyperarousal). The expected value of a stimulus A , Q^A , is updated as follows,

$$Q_{t+1}^A = Q_t^A + \alpha \cdot \kappa_t^A \cdot \delta_t \quad (2.3)$$

α represents the (fixed) learning rate, δ_t represents the prediction error at time t . The associability of stimulus A is represented by κ_t^A , updated as follows.

$$\kappa_{t+1}^A = (1 - \eta)\kappa_t^A + \eta|\delta_t| \quad (2.4)$$

The associability weight parameter, $0 \leq \eta \leq 1$, controls how much historic prediction errors influence the future associability of each cue.

The model was found to predict participant choices better than models without associability, suggesting that attention-based learning may be a key aspect in the underlying mechanisms of PTSD and may assist in refining treatment targeting. Similar to the momentum term used by Kaye et al., associability allows for an agent to take advantage of prediction errors gathered across its lifetime. In other words, both associability and

momentum terms add value to their respective models by allowing agents to utilise past experiences when generating future predictions.

Comparing these terms, momentum (Eq. 2.2) sums the decayed historic prediction errors of threat across all contexts to influence prediction, whilst maintaining a constant learning rate. Associability (Eq. 2.4) creates a more dynamic learning rate (seen in Eq. 2.3), where learning for more unreliable cues is updated preferentially. These terms provide us with an opportunity to explore the extent to which past experiences can influence future fear learning processes. In relation to cPTSD and childhood trauma specifically, investigating how these terms react to prolonged stressors in early life stages and how these influence the extinguishing of learned fear responses could prove to be key in developing the understanding of cPTSD. We propose a model combining associability with the momentum model of Kaye et al. in Chapter 4.

2.1.3 Ross et al. (2018) - General PTSD Learning Deficits

Ross et al. [23] explored how much PTSD is associated with general reinforcement learning, outside of the context of learned fear stimuli. The authors investigated variations in prediction errors between a group of PTSD individuals and a group of control individuals. Participants were shown two images of houses and were required to identify which was unlocked in order to maximise monetary reward. Results were fit to several adaptations of the classical Rescorla-Wagner model (Eq. (2.5)).

$$V_{t+1} = V_t + \delta * \alpha \quad (2.5)$$

The expected value of a choice is represented by V_t , the prediction error $\delta = (outcome_t - V_t)$, and α represents the learning rate ranging from 0 – 1. The authors found that a risk-sensitive, anti-correlated model fitted participant behaviour best. “risk-sensitive” implies separate learning rates are used for positive and negative prediction errors. “anti-correlated” implies that updates to the expected value of the unchosen stimulus are in the opposite direction of the prediction error. Choice estimate values are then updated via Eq. (2.6) and Eq. (2.7), where $\alpha^{+/-} = \alpha^+$ if $\delta \geq 0$ or α^- otherwise.

$$V_{t+1}^{Chosen} = V_t^{Chosen} + \alpha_{Chosen}^{+/-} \delta \quad (2.6)$$

$$V_{t+1}^{Unchosen} = V_t^{Unchosen} - \alpha_{Unchosen}^{+/-} \delta \quad (2.7)$$

Although no significant differences in parameter values across these adult groups were found, previous cognitive flexibility investigations using this model have been

performed on adolescents that found a significantly increased learning rate for negative prediction errors compared to adults [13]. Therefore, we propose that this model may have further implications (and provide more distinct results between groups) when applied to childhood trauma and cPTSD.

Comparing to previous models, this model is similar to the associability model of Brown et al., where associability can be portrayed as *weighting* learning such that unreliable cues are updated preferentially and influence learning more than others. Ross et al. do not include this weighted learning as such, although there are separate learning rates for positive and negative prediction errors, no specific cues are updated preferentially. The momentum model is also the only approach to explicitly incorporate non-associative learning (via the momentum term).

2.1.4 Yanamori et al. (2023) - Anxiety Related Approach-Avoidance

Yanamori et al. [31] investigated anxiety-related approach-avoidance, proposing that anxiety may increase avoidance responses (decrease approach responses) in an approach-avoidance task, which may be explained by increased sensitivity to punishments rather than rewards. In other words, anxious people are more likely to avoid pursuing reward than to pursue a reward associated with potential punishment. Although based on anxiety-related avoidance, the concepts used have parallels to PTSD-related avoidance.

Yamamori et al. conducted a “Restless Bandit” behavioural task where participants aimed to maximise rewards while avoiding punishments. Participants chose one of two images that resulted in either a reward or no reward. Options were grouped as “safe” and “conflict” (unknown to participants), where “conflict” options sometimes resulted in screaming sounds (punishment) and could be any combination of reward/no reward and punishment/no punishment (4 possible outcomes). “safe” options never resulted in punishment and were, on average, less likely to produce reward over “conflict” options. “conflict” options thus represented more reward but also the potential of punishment.

Participant choices were fit to variations of the Q-learning algorithm [28]. The model with specific learning rates and specific outcome sensitivity parameters for rewards and punishments fit behavioural data best. Probability estimates of outcomes are progressively updated as action outcomes are observed. These estimates are used by participants to choose an option that maximises subjective value, e.g. option most likely to result in reward, or option most likely to result in no punishment. The probability estimates of observing an outcome, $o = \{reward, punishment\}$, of a chosen option,

$a = \{conflict, safe\}$, were updated via the below:

$$Q_{i+1}^o(a) = Q_i^o(a) + \alpha \cdot [o^t - Q_i^o(a)] \quad (2.8)$$

Learning rate α is split into reward- and punishment-specific learning rates, α^r and α^p that imply a similar risk-sensitive approach to Ross et al., however, learning rates are chosen based on outcome, rather than sign of prediction error. Probability estimates are then merged into action weights, W , at every trial, with outcome sensitivity parameter, β , separated into reward- and punishment-specific parameters, β^r and β^p .

$$W = \beta^r \cdot Q^r - \beta^p \cdot Q^p \quad (2.9)$$

The action weights then form the basis of which choice is made between options, modelled with a softmax function,

$$P(a) = \frac{W(a)}{\sum_i W(i)} \quad (2.10)$$

The sensitivity parameters β^r, β^p capture the extent to which each outcome impacts choice and allow for asymmetries in how individuals value reward and punishment to be captured. The ‘‘reward-punishment sensitivity index’’, β^r/β^p , provides a unique index for each individual as to where they lie on the approach vs avoidance spectrum, higher values correspond to approach preference, and lower values to avoidance preference.

The authors found that punishment learning rate α^p and reward-punishment sensitivity index β^r/β^p were negatively correlated with task-induced anxiety. This indicates that anxious individuals were slower in updating estimates of punishment probability, and that they placed more weight on punishment relative to reward when choosing between options (reward was biased by potential for punishment). This explains the effect of task-induced anxiety on avoidance behaviours, and may be useful in providing further insight into mechanisms that drive avoidance. As anxiety disorders are often comorbid with PTSD [6], similar experiments could be conducted on PTSD individuals exploring if parameter values for PTSD individuals are similar to those in this anxiety-related study. Differences in parameter values may provide insight on differences between the mechanisms of these disorders and provide potential explanation for comorbidity.

2.1.5 Eldar et al. (2016) - Momentum Model of Mood in Bipolar Disorder

Here we detail the momentum model by Eldar et al. [10], which is implemented by Kaye et al. in the main model of this dissertation. This model aims to describe mood in

a computational model of bipolar disorder (BD). It is known that BD can stem from prolonged childhood trauma (cPTSD) and individuals often suffer from both disorders [6]. Also, cPTSD is often misdiagnosed as BD due to similar symptoms, as such we can draw parallels to cPTSD mechanisms.

Eldar et al. propose that experiences affect mood, which in turn influences future experiences. They suggest that mood is the overall momentum of recently experienced outcomes, and that its biased influence on outcome perception allows for the “correction” of learning, accounting for environmental dependencies. The model is based on an adjusted form of the Rescorla-Wagner model (Eq. (2.5)), this standard form assumes different states are independent of each other. However, Eldar et al. suggest that if different states are *not* independent of each other, and multiple states can be similarly affected by some environmental factor, then an adjustment is required that updates expectations of all states affected by this factor, when a prediction error is experienced in any one of the states. The *momentum* term is introduced to represent this:

$$m_{t+1} = m_t + \alpha((outcome_t - V_t^s) - m_t) \quad (2.11)$$

Momentum assumes that states close in space and time are affected similarly. The term is combined with Eq. (2.5) to create the momentum model for mood (Eq. 2.12), where f_t is a scaling factor. This model allows for expectation of reward in a given state to reflect outcomes experienced in all states.

$$v_{t+1}^s = v_t^s + \alpha(f_t * m_t + (outcome_t - v_t^s)) \quad (2.12)$$

The authors apply behavioural data (gathered by Eldar and Niv [9]) from BD individuals who played a different slot machine before and after a wheel-of-fortune draw that they either won or lost. Participants who reported high emotional instability were asked to report how their mood varied throughout the experiment. The wheel-of-fortune outcome was shown to influence mood in these individuals. Those who won the wheel-of-fortune preferred the second slot machine (after the win, played in a better mood), while those who lost the wheel-of-fortune preferred the first slot machine (before the loss, played while in a better mood), showing how outcomes that may not be directly related to future stimuli can influence interpretation of future stimuli (i.e. mood influences learning).

Kaye et al. implement this concept of mood into a model investigating the fear learning aspect of PTSD, whereby they propose that trauma experienced in any context should hold some influence over threat predictions in other contexts that are close in space and time. We now move on to the evaluation of this main paper.

Chapter 3

Re-implementation & Evaluation of Kaye et al. (2023)

The majority of previous RL models of PTSD focus on associative learning, describing the disorder as a deficiency in extinction learning of conditioned fear responses thought to rely on RL mechanisms, i.e. future predictions are influenced by errors between predictions and actual outcomes (prediction errors). Many studies build upon the classical Rescorla-Wagner model [21] which accurately describes associative learning. However, Kaye et al. note that these studies do not incorporate non-associative learning, and thus fail to account for how repeated trauma (across different contexts) can influence future threat predictions. There is currently a lack of knowledge surrounding how non-associative learning in PTSD can be described computationally, as well as how this process may be implemented neurobiologically.

The authors propose a novel RL model to advance this knowledge, combining the basic Temporal Difference (TD) RL model and the mood model by Eldar et al. [10] (section 2.1.5). Kaye et al. suggest that the estimation of *frequency* of trauma may be influential to PTSD adaptations in the brain, instead of simply threat associations with specific cues. The authors first propose a Bayesian model for estimating frequency of threat, to understand how well an agent can perform based on its own experiences. Two RL models are then created; the simpler TD model, and the TD-Momentum model (our main focus). These RL models are fit to data from a stress-enhanced fear-learning (SEFL, section 3.2) model of rodents to assess if the TD-Momentum model provides an improved fit to threat learning.

3.1 Bayesian Baseline Model

Kaye et al. propose that an agent in an environment must estimate future likelihood of trauma based upon previous experiences, which can be defined in a Bayesian process. This Bayesian model is used as a baseline for threat estimation from the perspective of an “ideal observer” that utilises all information of a stimulus (in the given context) to estimate threat as best as possible. This allows us to measure how much information an agent can accumulate over a lifetime when subjected to traumatic events.

A simple probabilistic model was used to test the performance of an ideal Bayesian observer in an information-poor environment, i.e. where agents have no previous knowledge of the environment and must learn from scratch. The model runs for 700 timesteps where, at each step, attacks occur with probability p_a and death occurs with probability p_d (given an attack occurs). This procedure is repeated for all timesteps, or until an attack results in death. If an agent survives a timestep, it continues to learn and this procedure repeats at the next timestep (model structure is shown in Appendix A, Fig. A.1). Agents have no power over the environment and do not know the fixed probabilities of attack (p_a) or death (p_d), these are set by researchers. During simulations, if the actual probability of death is set too high, agents will not gain enough experience over their lifetime in order to produce good estimates. Therefore, agents must maximise the information available in order to produce accurate estimates.

3.1.1 Bayesian Attack Rate Estimation Model

In the attack model for Bayesian estimation, attacks are binomially distributed over the agent lifetimes (700 timesteps), with a probability of attack $p_a = 0.2$ and a probability of death given attack $p_d = 0.01$. Both p_a and p_d agent estimates are calculated via Bayes theorem, based on a sequence of attacks, \mathbf{x}_t .

$$p(p_a, p_d | \mathbf{x}_t) = \frac{p(\mathbf{x}_t | p_a, p_d) p(p_a, p_d)}{\int p(\mathbf{x}_t, p_a, p_d) d\mathbf{x}_t} \quad (3.1)$$

Upon new observations, posterior estimates are updated through an affine invariant Markov Chain Monte Carlo (MCMC) sampler [11]. As time progresses, the previously evaluated posterior becomes the new prior and estimation is repeated. As agents have no starting knowledge, we set a flat initial prior at $t = 0$. The MCMC sampler was fit using the below likelihood function to estimate the posteriors of p_a and p_d .

$$\ln(\mathcal{L}(\mathbf{x}_t, p_a, p_d)) = \sum_{attacks} \ln(p_a(1 - p_d)) + \sum_{non-attacks} \ln(1 - p_a) + \sum_{deaths} \ln(p_a * p_d) \quad (3.2)$$

The affine invariant MCMC sampler, first proposed by Goodman and Weare [12], is applied due to its efficiency in exploring skewed posterior distributions. Here, agent posterior estimates are predicted to be positively skewed between 0 and 1, with more estimations being near the real value of p_a . This estimator generates the best estimate available to an ideal Bayesian observer, with the information available at each timestep

3.1.1.1 Bayesian Attack Rate Estimation Model - Implementation of Methods

In terms of implementing models, all code was written in Python (Version 3.9 [27]) using various packages. Attack sequences of 700 timesteps based on figure A.1 were simulated for use in the MCMC Sampler, code for these simulations is given in Appendix B.1. The “*EnsembleSampler*” function from the “*emcee*” [11] package was used to generate the MCMC sampler with 30 walkers, “*num_steps*”=60, “*burnin*” = 0.3 and a “*moves*” argument of “*Stretch_move*” with stretch parameter = 2. Code for the MCMC implementation and related plots can be found in Appendix B.2 - B.3.

Kernel Density Estimation (KDE) was applied to results of the MCMC sampler using the “*gaussian_kde*” function of the “*stats*” package [1]. Suitable bandwidths were selected for the p_a and p_d distributions to generate smooth posteriors. A “*gaussian_filter*” was applied so that contour plots of p_a against p_d were readable.

3.1.2 Auto-regressive Time Series

Auto-correlated attack rate time series were also generated to create attack sequences used as input to the Bayesian model, following the auto-regressive (AR) process below.

$$p_{a,t} = cp_{a,t-1} + \mathcal{N}(0, 0.01) \quad (3.3)$$

Attack rate at timestep t is denoted $p_{a,t}$, the correlation for successive timesteps is represented by c , additional noise $\mathcal{N}(0, 0.01)$ is sampled from a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 0.01$. The resulting time series of unique p_a at each timestep is used to generate attack and death sequences for input to the MCMC sampler. Whilst p_a varies at each timestep, $p_d = 0.01$ remains constant. This process generates attack sequences where $p_a \approx 0.01$, as $p_{a,0} = 0$ (initial p_a).

3.1.2.1 Auto-regressive Time Series - Implementation of Methods

AR time series simulations used the “*ArmaProcess*” function of the “*statsmodels.tsa*” package [25]. This function requires “*ar*” and “*ma*” arguments which represent

coefficients for the auto-regressive and moving-average lag polynomials, respectively (zeroth lag always set to 1). To implement the AR process in Eq. (3.3), arguments are set to $ar = [1, -c]$ and $ma = [1]$, where c values must be negated in the ar argument, e.g. for $c = 0.7$ we set $ar = [1, -0.7]$. As we require a strictly non-negative probability of attack, p_a , in the binomial distribution, time series must be clipped between 0 and 1 before generating attack sequences (related code detailed in Appendix B.5).

We generated 10,000 separate lifetime (700 timesteps) attack rate time series, used to create 10,000 separate attack sequences consisting of 0's (no attack) and 1's (attack). Each sequence was used as input to the MCMC sampler, testing how the model reacts to varying AR correlation coefficients c (related code in Appendix B.6).

3.1.3 Bayesian Model - Re-implementation

Here we re-implement this Bayesian model and reproduce plots similar to those of Kaye et al.. Fig. 3.1 shows posteriors generated from MCMC sampler results of a typical sequence of attacks over an agent lifetime, created using a single AR time series of attack probability (correlation constant $c = 0.7$). This shows estimates of the probability of attack (p_a) to be very accurate, with real attack rates generated by the AR time series being $p_a \approx 0.01$, our p_a posterior estimates are clearly centred around this value, ranging between 0 - 0.02, indicating low variance. The accuracy of p_d posterior estimates is very poor in comparison, with estimates over a much wider range, showing non-convergence to a confident estimate and indicating the model is poor at estimating lethality of attacks. This is intuitive when considering the model structure; an agent cannot provide further information for posterior updates after death (Fig. A.1). This Bayesian model progressively improves its estimates over time, becoming more confident in estimating attack frequency throughout its life (seen in Appendix A Fig. A.2, variance in model estimates decreases over time, showing convergence for p_a estimates). The correlation in AR time series can influence precision of Bayesian model attack estimates, larger auto-correlation ($c = 0.999$) results in larger variance in p_a estimates (flatter, wider posteriors) and less accuracy in results. Lower auto-correlation ($c = 0.7$) results in lower variance in p_a estimates (sharper, thinner posteriors) and more accuracy in results (seen in Appendix A, Fig. A.3 showing “average” MCMC sampled densities for 10,000 simulated AR time series with various values of correlation).

Kaye et al. then briefly discuss links to childhood trauma, comparing an early life stress (ELS) scenario to one of lifetime stress. We re-create a similar plot in Fig. 3.2,

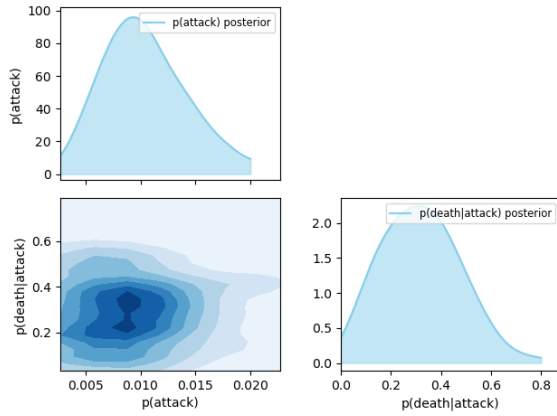


Figure 3.1: Estimated posteriors for p_a and p_d for typical attack sequence generated from AR time series of attack rates with $c = 0.7$. The Bayesian estimator approximates p_a very accurately ($p_a \approx 0.01$), however estimates lethality of attacks poorly, as agents gain no information after death.

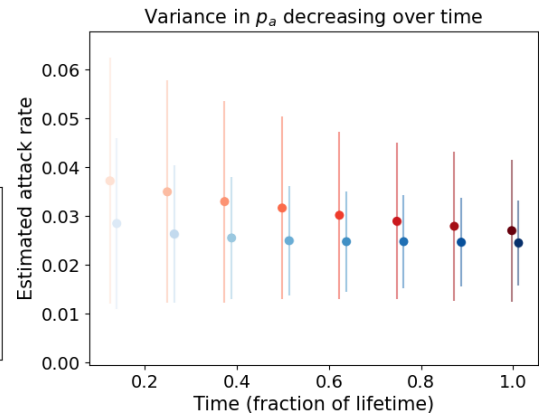


Figure 3.2: ELS scenario (red) with $n = 15$ attacks in early life and lifetime scenario (blue) of $n = 15$ attacks across life. The ELS scenario shows larger threat predictions and increased variance across its lifetime, showing the Bayesian model can naturally represent the impact of ELS.

where $n = 15$ attacks are uniformly distributed in the first half of life to represent ELS, and across the entire lifetime for the lifetime scenario. Attacks in early life can cause mean and overall variance of threat estimations to be larger (red) compared to attacks spread across a lifetime (blue), showing how the Bayesian model can naturally represent the disproportionate impact of early life trauma (see Appendix B.7 for related code).

3.1.4 Bayesian Model - Additional Investigations

We now extend upon the findings of Kaye et al. by applying our own ELS scenarios and analysing how the model reacts. We investigated the effects of ELS clusters of 3 attacks compared to random attacks over a lifetime (see Appendix B.8 for related code). For the ELS scenario, 5 equally spaced ELS clusters of 3 attacks were administered (over adjacent timesteps, totalling 15 attacks), while the lifetime scenario consisted of 5 attacks randomly over a lifetime. Clustered attacks represent more significant or prolonged traumatic ELS compared to single attacks. Fig. 3.3 shows the ELS scenario has larger mean and variance of threat estimations compared to the lifetime scenario.

However, this may be due to the ELS scenario simply having more attacks. We also compare the clustered ELS scenario to a lifetime scenario with the same number

of attacks. Fig. 3.4 shows mean threat predictions in the lifetime scenario all increase slightly as expected. However, variance in lifetime threat estimation remains similar to lifetime variance in Fig. 3.3. The mean threat estimations in the ELS scenario, as well as variance, still remain significantly larger than those in the lifetime scenario.

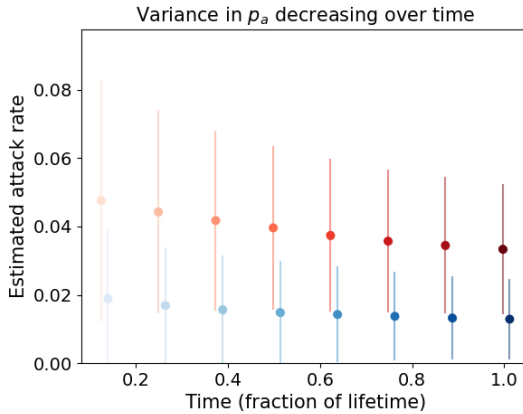


Figure 3.3: 5 equally spaced 3-clustered attacks (red-ELS), compared to 5 random attacks over lifetime (blue-lifetime)

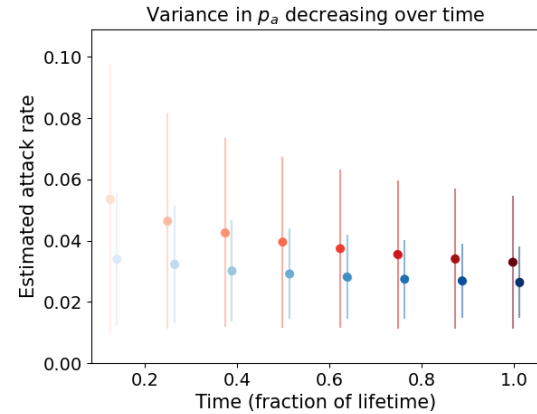


Figure 3.4: 5 equally spaced 3-clustered attacks (red-ELS), compared to 15 random attacks in lifetime (blue-lifetime)

Comparing the clustered ELS scenarios in Figs. 3.3 and 3.4 to random single ELS attacks in Fig. 3.2 (all red plots), we observe that the clustered scenario threat estimates maintain larger variance and mean in comparison to the single attack scenario. In other words, these simulations show that clustered attacks in early life can cause the Bayesian estimator to be less accurate and predict higher threat than random single attacks in early life. However, as these plots show single simulations, they may not generalise to every scenario of clustered vs single attacks. In any case, there is a clear disproportionate effect on threat estimates caused by any early life attacks, clustered or single.

3.1.4.1 Bayesian Model Implications for cPTSD

In terms of cPTSD implications, our investigations correspond to comparing types of early life trauma to trauma experienced over a lifetime. Findings suggest that individuals exposed to clusters of ELS predict higher threat across their lifetime compared to those who suffered less prolonged ELS, or less intense trauma spread across their lifetime. These increased threat predictions represent an increased perception of threat in the world throughout day-to-day life which may provide explanation for the dysfunctional physical and emotional responses prevalent in cPTSD, classified by symptom clusters.

Physical symptoms may be represented by an increased state of hyperarousal (startle response) and hypervigilance to sensory inputs. Emotional symptoms would be represented by emotional dysregulation during various situations. For example, increased emotional reactivity to small stressors, where an individual may react with intense negative emotion to a minor stressor (e.g. being easily overwhelmed by a minor inconvenience). Another possible representation may be a persistent, deep sense of mistrust of the world in general, caused by persistent elevated threat level. Symptoms may then lead to impairment in important areas of social functioning (occupational, family, educational) and may also be linked to difficulties maintaining relationships.

It could also be argued that this increased perception of threat may lead to avoidance of reminders of the trauma itself, or even to avoidance of small stressors unrelated to the original trauma that, when encountered, may trigger these aforementioned symptoms.

3.2 Stress-Enhanced Fear-Learning Model

Here we define the Stress-Enhanced Fear-Learning (SEFL) model of rodents which produces behavioural data used for fitting to the RL models of threat. The SEFL has been shown to result in long-lasting enhancement of fear behaviour in rodents [20]. Stressed mice were subject to 15 unpredictable footshocks in context A on day 1 over 90 minutes, with random inter-shock intervals of 4 to 8 minutes. Mice were transferred to context B (with completely different sensory characteristics) for 10 minutes on day 6 and administered a single footshock after 5 minutes. On day 7, mice were returned to context B for a further 10 minutes with no footshocks. A group of control mice were not exposed to any footshocks (unstressed) on day 1 in context A, but had the same experience as stressed mice on days 6 and 7. Freezing was defined to be complete cessation of movement (other than breathing), and was measured for all mice across all days. Stressed mice showed an increased average freezing percentage across days 6 and 7 compared to unstressed mice, shown in Figs. A.4 and A.5 of Appendix A.1.

“Freezing” is a voluntary defence mechanism in rodents where the animal is completely immobile in response to potential threat. The raw freezing time series created for each mouse is binary, where 1 marks observed freezing behaviour and 0 marks no freezing. This is transformed into “smoothed freezing”, used as input for the models. Smoothing is a pre-processing technique to assist in time series analysis of freezing in which the percentage of freezing within consecutive time intervals is calculated, giving a smoother view of freezing variation over time. Kaye et al. use a 15s sliding window

smoothing, i.e. average freezing is calculated over 15s fixed intervals, the window is then moved on 1s and freezing percentage is calculated again, creating a time series of moving average values between 0 and 1 (an example of this smoothed freezing is shown in Fig. 3.7). We perform a similar smoothing to the raw freezing data, generating our own smoothed freezing sequences (related code can be found in Appendix C.2).

3.3 Reinforcement Learning (RL) Models

We now introduce the RL models applied by Kaye et al. that aim to explain the neural processes of fear learning in PTSD. The free parameters in each model allow for fitting to the stressed and control (unstressed) mice groups from the SEFL, allowing for hypotheses and conclusions to be drawn from comparisons across groups. We begin by re-implementing the models and reproducing similar plots before performing additional investigations on TD-Momentum model behaviour and response to ELS.

Fitting the TD and TD-Momentum models to the contextual SEFL experiment data, Kaye et al. tested if threat momentum was a source of PTSD symptoms, rather than just the specific association with traumatic events. Parameters for both models were fit via maximum likelihood, minimising the negative log likelihood (NLL) function of each model. The maximum likelihood fit for each model (of each mouse) was compared using BIC scores to identify the best fitting model for each mouse. Inputs to models were binary time series of attack sequences (1s indicating attacks, 0s no attack), threat was fit for both contexts across all three days, where the threat variable in models (output) was re-scaled to (0.1, 0.9) to match the real smoothed freezing probabilities.

3.3.1 Parameter Recovery

Parameter recovery was performed for both models prior to fitting. Generative model functions were built and two scenarios for parameter sampling were carried out: one set from multivariate normal distributions (with no covariance), and one set from random choices of potential parameter value in lists provided by Kaye et al..

The multivariate normal distribution had means equal to the midpoints of each respective parameter value list, with suitable variances to ensure sampled values were spread across the complete range of values for each parameter. Parameter values sampled from the lists were simply random choices from each list. 100 simulations were run for each sampling scenario and parameters were recovered via maximum likelihood fit,

minimising the NLL function of each model using the “*differential_evolution*” function with a “*polish = True*” parameter from the “*scipy.optimize*” package [15]. The inclusion of “*polish*” means the “*L-BFGS-B*” method is applied to the best population member, slightly improving results. Differential evolution was found to be the best optimisation method for both models. Related code can be found in Appendices C.4-C.5 (TD model), and Appendix D.4-D.5 (TD-Momentum model).

Recovery was very successful for the basic TD model, with Pearson correlation coefficients between sampled and recovered parameter values for both α and γ_1 parameters being close to $r_\alpha \approx r_{\gamma_1} \approx 1$ (perfect recovery), showing the TD model function and related TD NLL function were implemented correctly (Figs. A.7, A.6 in Appendix A). Recovery for the TD-Momentum model was successful for α with correlation consistently returning values $r_\alpha \approx 1$. Recovery for γ_1 was not as successful, with correlation in the range $0.57 < r_{\gamma_1} < 0.87$. Recovery for γ_2 and f parameters was less successful with correlations coefficients in the ranges $0.38 < r_{\gamma_2} < 0.69$ and $0.2 < r_f < 0.61$ (Figs. A.8, A.9 in Appendix A), this weaker recovery may be due to the parameters having a similar effect on threat estimations, hindering the minimiser from accurately minimising the NLL, thus returning non-minimal values. No apparent correlation between fitted parameters was found, suggesting this poor recovery is not due to the parameters trading variance with each other. Kaye et al. do not note any parameter recovery in their study, it may be the case that recovery was not performed. In any case, we do not have any recovery figures or thresholds to compare our findings to. Overall, our findings show that both models and related NLL functions are correctly implemented.

3.3.2 TD Model

The Temporal Difference (TD) model, adjusted from Sutton and Barto [26], follows the update rule:

$$T_{c,t} = T_{c,t-1} + \alpha(u_t - \gamma_1 T_{c,t-1}) \quad (3.4)$$

Threat at time t in context c is denoted $T_{c,t}$, learning rate is represented by α and controls how much the prediction error ($u_t - \gamma_1 T_{c,t-1}$) updates future threat estimations, decay rate is represented by γ_1 and controls how much we discount the value of future rewards. Threat is learned from a sequence of unconditioned stimuli, u_t , as input, in this case a sequence of attacks (0 for no shock, 1 for shock). Parameter values are bounded as follows: $0.05 \leq \alpha \leq 0.9$ and $0.9 \leq \gamma_1 \leq 0.99999$. This update rule describes *associative* threat learning for PTSD.

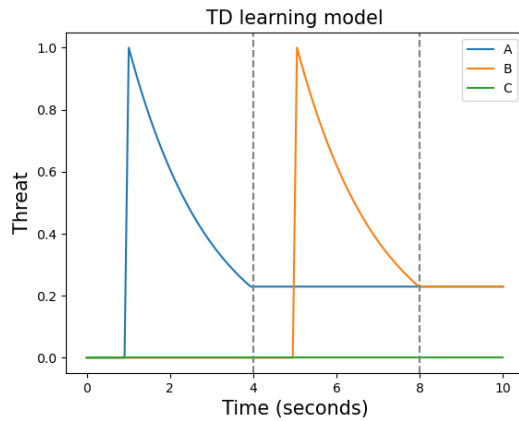


Figure 3.5: Typical TD threat estimation, single attack in context A and B, no attacks in C. Threat prediction for context C does not increase from 0. Dashed lines indicate moves from context A to B to C.




Figure 3.6: Typical TD-Momentum threat estimation for similar scenario. All threat estimations are affected when an attack occurs in any context. All threat in context C is due to attacks in other contexts.

3.3.3 TD Model - Re-implementation

The TD model allows us to model threat estimations for agents within the context in which attacks occur. This model assumes that threat prediction in each context is independent of other contexts, as such, there is no means for threat estimations to “carry across” contexts. Fig. 3.5 shows a simulation of a single attack in contexts A and B, with no attacks in C. Starting in context A, dashed lines represent transitions between contexts (A to B and B to C). There is no increase in contexts B and C threat when an attack occurs in A (the sharp increase in threat), similarly for A and C when an attack occurs in B. Predictions in any context remain constant upon leaving that context.

The TD model thus captures the *associative* fear learning of the disorder, allowing agents to associate traumatic experiences with the context in which they occurred. However, the simple TD model neglects the fact that an agent may have had previous experiences in other contexts that may influence future threat predictions across novel contexts. Kaye et al. argue that these past experiences influence an agent's current view of the world (or “mood” as referred to by Eldar et al. [10]), which needs to be incorporated for threat estimation to represent both associative and non-associative learning. They present the TD-Momentum model to express this.

3.3.4 TD-Momentum Model

The TD-Momentum model is an extension of the TD model with the addition of a “momentum” term, introduced by Eldar et al. [10] (section 2.1.5). The inclusion of momentum allows for prediction errors in any context to influence estimations in novel contexts, incorporating *non-associative* fear learning. The update rule is defined as:

$$T_{c,t} = T_{c,t-1} + \alpha(u_t - \gamma_1 T_{c,t-1}) + f m_t \quad (3.5)$$

$$m_t = m_{t-1} + \gamma_2 \sum_{c=\{A,B,\dots\}} \alpha(u_t - T_{c,t-1}) \quad (3.6)$$

Momentum at time t , m_t , is computed from the sum of decayed prediction errors across all contexts, with decay rate $0 \leq \gamma_2 \leq 0.8$. Scaling constant $0.01 \leq f \leq 3.0$ controls how much momentum influences threat prediction updates. Inputs are sequences of attacks and outputs are threat prediction sequences across all contexts.

3.3.5 TD-Momentum Model - Re-implementation

For the TD-Momentum model, threat estimations in unique contexts are no longer independent. Fig. 3.6 shows an example of an agent exposed to a similar simulation to Fig. 3.5. Comparing to the TD model in Fig. 3.5, attacks in any context influence estimations in all other contexts and Context C threat also increases, even though no attacks occur here; these effects are due to momentum. Estimations also plateau slowly when leaving a context, rather than abruptly plateauing like the TD model.

3.3.5.1 Re-implementation - Fitting to SEFL Data

We now reproduce the fitting and model comparison performed by Kaye et al. An example of fitting the smoothed freezing to the TD-Momentum model is shown in Fig. 3.7 (single stressed mouse ID G:34). This shows smoothed freezing behaviour observed on each day (top row), as well as the related TD-Momentum fitted threat (bottom row). The TD-Momentum model fits threat well, and shows the disproportionate freezing response of stressed mice to a single footshock on day 6 in context B. Kaye et al. note that this sensitised behaviour on day 6 is explained by the momentum term, which allows threat estimates to be linked across context A on day 1 and context B on day 6. For unstressed mice, TD and TD-momentum fits were very similar due to these mice not experiencing any shocks on day 1, thus not showing the same level of sensitised threat to the single shock on day 6 seen in stressed mice.




Figure 3.7: Observed smoothed freezing of stressed mouse (ID: G34) (top row) and corresponding TD-Momentum fitted threat (bottom row) across all days of the experiment. Blue and orange lines represent threat estimates in contexts A and B, respectively. Vertical lines represent input (footshocks) in contexts A (green) and B (red).

Fig. 3.8 shows the model comparison for all mice. Bayes Information Criterion (BIC) scores for maximum likelihood fits of each model were computed for each mouse and the difference ($BIC_{TD} - BIC_{TD\ Momentum}$) was taken to evaluate the preferred model. Stressed mice favoured the TD-Momentum model (14/15 mice), and control (unstressed) mice favoured the basic TD model (16/18 mice).

Taking the approach that the momentum model may be favoured as it allows for influence across contexts, we could argue that unstressed mice may have previously experienced threatening events in other contexts (before the experiment) that influence the freezing responses observed in the SEFL study. As such, it would be intuitive to hypothesise that *all* mice should have favoured the TD-Momentum model.

This leads to the question of how long momentum should persist after a series of events. Kaye et al. found this to depend on how long trauma itself persists. When attacks are correlated in time, including momentum provided a better estimation of true threat level in an environment. When attacks are uncorrelated, they found that momentum holds no advantage and the optimal $f = 0$ reduces to the basic TD model (Fig. A.10 Appendix A). As unstressed mice experience uncorrelated attacks (no attacks) on day 1, the f parameter tends towards 0, making the two models equivalent. As the TD-Momentum model has 4 parameters and the TD model has 2, the BIC




Figure 3.8: Model comparison for 15 stressed (red) and 18 unstressed (blue) mice. BIC scores were computed from the maximum likelihood fits of both models for each mouse.

scores for unstressed mice penalise the TD-momentum model more heavily, leading to the majority of control mice favouring the TD model. Stressed mice experienced correlated attacks on day 1, including the momentum term predicted increased freezing in the novel context (B) which resulted in a better fit, with momentum accounting for improved predictions compared to the TD model (related code in Appendices C.2-C.3 & D.2-D.3).

3.3.6 TD-Momentum Model - Additional Investigations

3.3.6.1 Additional Investigations - Model Behaviour

We now extend upon the findings of Kaye et al. with an analysis of TD-Momentum model behaviour, focusing on the additional parameters of momentum decay rate (γ_2) and scaling (f). Low values for these parameters encourage a slow summation of prediction errors across contexts and a normal observation of threat estimates summing to influence other contexts. However, larger values result in oscillatory behaviour of threat estimations across all contexts, with γ_2 being largely responsible for this as it controls the decay rate of summed prediction errors within the momentum term.

Fig. 3.9 shows TD-Momentum threat in response to 6 random ELS attacks. Typical values of $f = 0.2, \alpha = 0.07, \gamma_1 = 0.9999$ were used with the upper bound value of $\gamma_2 = 0.8$. This large γ_2 value causes context B threat to be larger in general. As momentum decreases after exiting context A, threat in both contexts reduces, threat




Figure 3.9: Maximum value of $\gamma_2 = 0.8$ causes oscillatory behaviour in threat estimations across all contexts. Attacks only occur in context A (blue highlight).

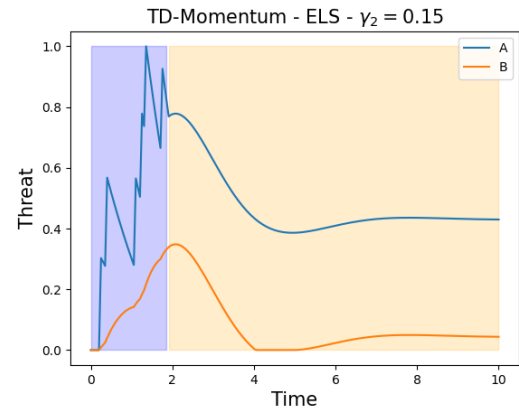


Figure 3.10: Similar to Fig 3.9, with $\gamma_2 = 0.15$. Small values of γ_2 can cause threat predictions to increase after reaching 0 for context B, even with no attacks.

in context B returns to 0 before a sudden increase into oscillatory behaviour. Fig. 3.10 shows that increasing γ_2 to as small as $\gamma_2 = 0.15$ can cause strange behaviour in estimations, where threat increases slightly after context B threat has reached zero. Maximising both parameters $f = 3$ and $\gamma_2 = 0.8$ results in extreme oscillatory behaviour for threat estimation which, in real life, would correspond to individuals experiencing cyclical periods of heightened sense of threat followed by periods of low sense of threat.

3.3.6.2 Additional Investigations - ELS Scenarios

Here we extend the findings of Kaye et al. by investigating the effects of repeated ELS on the TD momentum model. This has links to extinction learning (the basis for exposure therapy); the process by which associative threat responses are reduced by re-exposure to the original trauma context with no threat stimulus. This produces small, negative prediction errors that reduce threat in the trauma context over time.

Fig. 3.11 shows an ELS scenario of 6 random attacks in context A (none in context B). If the original context of repeated trauma is only visited briefly (e.g. for the duration of the trauma), after which it is exited, threat in both contexts plateaus at high level, explaining sensitisation (increase in response) to repeated threat. Threat level will not decrease further until the original trauma context is returned to and extinction learning can re-start, reducing threat in all contexts (via reduced momentum). Fig. 3.12 shows how threat estimations change when the original trauma context is re-entered, allowing for extinction learning to begin again. Threat for both contexts reduces more rapidly

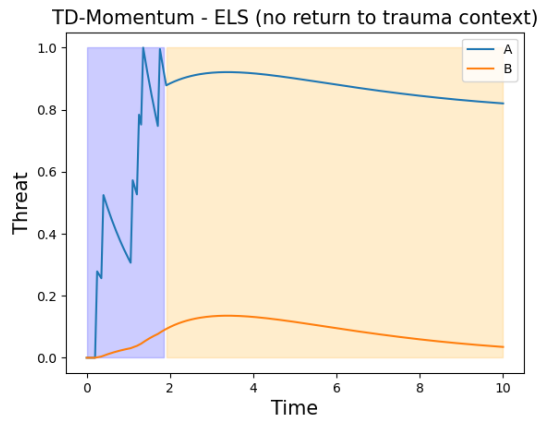


Figure 3.11: ELS of 6 random attacks in context A influences momentum and thus threat estimations in both contexts (blue for A, orange for B). Moving from context A to B, threat estimations both plateau.

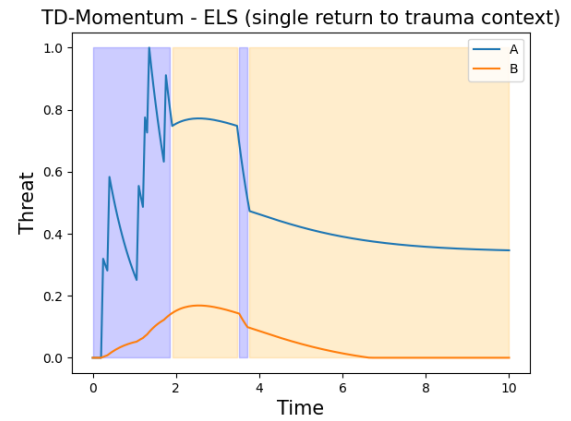


Figure 3.12: Similar to figure 3.11, however here we re-enter context A briefly before moving back to context B. Threat in both contexts decreases via momentum due to re-visiting trauma context A.

while in the original trauma context, even a short return to the original trauma context can have a large impact due to the larger prediction errors influencing momentum, context B threat actually returns to 0 following the re-entry. The rate of threat reduction depends on parameter values; larger learning rates, α , result in more rapid decrease.

We also investigate potential reasons for habituation, the phenomenon in which threat response decreases when exposed to repeated trauma. This has been shown to have links to numbing symptoms seen in cPTSD [17]. We compare ELS clusters followed by smaller singular threat with a scenario of singular, evenly spaced ELS. Fig. 3.13 shows threat estimates for an agent exposed to 5 singular, evenly spaced ELS instances in context A (none in context B). Threat in context A increases to around the same level after each attack, threat in context B slowly increases after each attack due to momentum. Upon entering context B, both threat levels reduce towards 0.

Fig. 3.14 shows the clustered ELS scenario where the first attack is replaced by a cluster of 4 attacks close in time, these represent more intense trauma in very early life. This cluster is followed by 4 single evenly spaced attacks. The initial cluster has a large impact on overall future threat estimations, whereby it creates a bias towards such intense threat and single attacks then result in a lower threat estimation in comparison to Fig. 3.13. This decrease in threat for prolonged trauma shows how habituation can occur in the model due to intensely clustered ELS. Momentum in the clustered ELS scenario influences context B threat largely, increasing in response to the cluster but

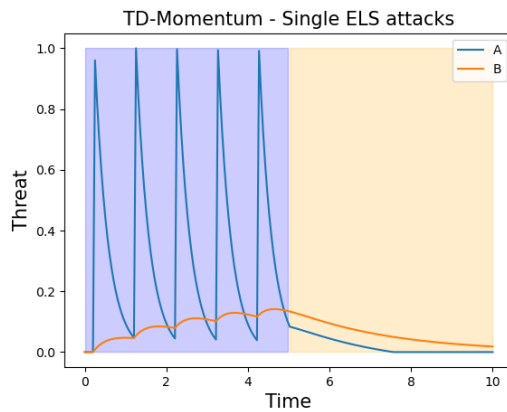


Figure 3.13: Single ELS scenario results in threat estimation in both contexts increasing with every attack (sudden spike).

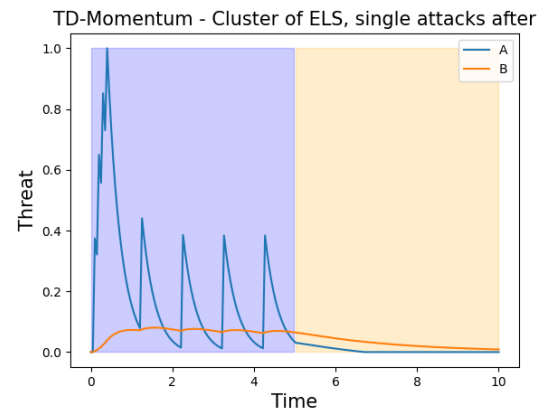


Figure 3.14: Clustered ELS scenario shows decreased response to single attacks, representing habituation in both contexts.

then slowly decreasing overall in response to singular attacks. This contrasts greatly to the single ELS scenario, where context B threat continually increases with every attack.

3.3.6.3 Additional Investigations - Implications for cPTSD

Here we review our previous findings with regards to implications for cPTSD and childhood trauma. Reviewing the oscillatory behaviour (Figs 3.9 and 3.10) related to variations in the γ_2 parameter, low γ_2 values appear “healthy”, providing a natural summing of all historic contextual prediction errors, and a good account of non-associative threat learning. However, the oscillatory behaviour caused by large γ_2 values may be representative of fear learning maladaptations that result in issues with affect regulation and emotional reactivity typical of cPTSD, leading to heightened emotional and physical responses to small stressors. If this oscillatory behaviour began at an early age due to ELS, this would naturally lead to persistent disturbances in many aspects of life (family, social, educational). The duration of momentum persistence would also impact the timescale of such symptoms.

Oscillatory behaviour of threat predictions may also explain failures in extinction learning, and thus provides explanation for the large percentage of treatment dropout and failure for exposure therapy. When behaviour is on a downwards “recovery” slope, where extinction learning is functioning correctly, exposure to small stressors may be successful in reducing the learned threat response and the patient may seem to be recovering. However, when an oscillation occurs causing an upward “relapse” slope, exposure to small stressors may amplify and result in sensitisation (a drastically

increased threat level), causing a failure in treatment and a persistent sense of heightened fear even after therapy. Individuals may then also avoid future therapy due to sustained elevated fear.

Regarding the ELS scenarios in Figs. 3.11 and 3.12, these provide potential explanations for the symptom clusters of avoidance and alterations in arousal and reactivity. These figures show threat to plateau at high levels due to momentum if the original trauma context is not returned to. This persistent elevated sense of threat may influence avoidance behaviours, i.e. contexts and situations may be avoided due to them being reminiscent of trauma. The remaining momentum due to the original trauma context then cannot be reduced as individuals avoid the context itself as well as reminders, meaning that threat remains at a fixed (albeit different) level for all contexts. Repeated, prolonged trauma across various contexts would amplify this effect, where individuals may avoid reminders of all original trauma contexts and thus avoid any opportunity for the combined momentum to reduce. This could be particularly prevalent in victims of childhood trauma and cPTSD, due to these individuals experiencing more prolonged trauma. Indeed, for those who experienced abuse from a trusted person (e.g. a family member) where abuse frequently took place in a childhood home, avoidance of reminders of this place would be an example of how momentum is unable to decrease. Thus causing a higher sense of threat in contexts with little association to the original trauma context, prompting heightened emotional (mistrust) and physical (hyperarousal) responses to events that may not be related to the original trauma.

Finally, we saw in Figs. 3.13 and 3.14 how habituation can occur in the model due to intensely clustered ELS. We noted how threat in context B was maintained at a particularly low level (via momentum), slightly decreasing as time moved on even though trauma was still occurring repeatedly. Threat in non-trauma contexts remaining at a low level (even decreasing) indicates that even though events are still occurring which may be extremely traumatic, the individuals perception of such events is no longer classing them as threatening, and so they are not affecting threat predictions in other contexts as they would in a “healthy” individual. In other words, this unresponsive momentum term could be viewed as the individual becoming accustomed to such trauma, habituating their responses and effectively numbing them to such traumatic events in all contexts, leading to the emotional numbing symptoms typical of cPTSD.

3.3.6.4 Additional Investigations - Conclusion

To summarise, the TD-Momentum model presented by Kaye et al. incorporates the associative and non-associative aspects of fear learning in PTSD, where momentum allows for threat in any context to influence threat prediction across all other contexts.

Various ELS scenarios were investigated, and we proposed that the TD-Momentum model can be representative of various dysfunctions in fear learning which may be reason for several symptom clusters in cPTSD. We also investigated the oscillatory behaviour caused by momentum with larger values of the γ_2 parameter, noting that we may expect “healthy” individuals to be associated with lower values, and more symptomatic, “unhealthy” individuals to be associated with larger values. A moderately influential, “healthy” momentum term would represent appropriate non-associative fear learning across various contexts. However, highly influential, “unhealthy” momentum biases fear learning, causing oscillatory behaviour which may be linked to issues with affect regulation in cPTSD. Additional research into how this parameter varies across control and disorder groups may provide further explanations. If differences in the γ_2 parameter (or any other parameter for that matter) are found, this could provide avenues for the application of this model as a tool for determining disorder risk and/or trajectory.

In terms of creating a human behavioural task to investigate this (where results are fit to the TD-Momentum model), it’s important to note that Kaye et al. define a single value function that simply estimates threat and has no “choice” to make. This could be adapted however, the value function of threat prediction could be applied to choices made in a human behavioural experiment where a task may be to learn threatening images over the course of several days (e.g. threatening images mixed among non-threatening images, varying in vagueness/difficulty). TD-Momentum fits could then be computed and differences between subject groups (control vs PTSD) compared to provide potential distinctions between levels of disorder severity, or even PTSD/cPTSD. Following the success of previous studies [7][14][23], combining behavioural data with neuroimaging may help to elucidate links between the model and how the neuromodulatory implementation of processes is represented in the brain.

Kaye et al. note the binary expression of attack and threat estimation as a drawback (0 for no shock/no freezing, 1 for shock/freezing). Obviously for humans the perception of attacks and threat is not so simple; both can be based over a broad spectrum that may vary across individuals. The authors propose that more precise manipulations of threat prediction errors over time may improve model validation.

Chapter 4

Extending the TD-Momentum Model of Kaye et al. (2023)

In this chapter we propose extensions to the TD-Momentum model that add value and explainability to improve the understanding of the disorder. We take several different approaches, noting motivation and advantages/disadvantages for each.

4.1 Incorporating Associability in TD-Momentum Model

Here we investigate the properties of the associability term proposed by Brown et al. [7] (section 2.1.2), and incorporate this into the TD-Momentum model.

Both the TD and TD-Momentum model base learning on updates via prediction errors; agents learn when a prediction is different to the actual outcome. Brown et al. incorporate this prediction error based learning with a dynamic associability term, assigning trial-by-trial associability values to each cue that scale over time with the historic prediction errors associated with that cue. Cues with a history of being unreliable/surprising receive larger associability which increases the dynamic learning rate related to that cue, causing a larger response. Thus, learning in this model is gated by the trial-by-trial associability value of each cue.

Incorporating this into the TD-Momentum model for threat prediction, where specific cues were given associability values by Brown et al., we propose that separate contexts are given associability values. The associability value of each context represents its influence to learning, and scales with the history of prediction errors that have occurred within that context. Therefore, contexts in which more traumatic events occur result in larger prediction errors and receive larger associability. Here, associability can

be viewed as a context-dependent momentum, that only (directly) influences learning within its related context.

Brown et al. define associability to represent the attention given to a cue over time, they found that learning in individuals with PTSD was modulated by this attention-based learning, as they exhibited an increased learning response (increased learning *rate*) to unexpected cues. We hypothesise that such individuals may have similar sensitised learning responses to high associability contexts (highly traumatic contexts).

4.1.1 Associability TD-Momentum Model Form

Combining Eq. (2.1) and Eq. (2.3), the proposed model takes the form below. Momentum is updated identically to the standard TD-Momentum model, and the trial-by-trial associability value of context c at time t is denoted $\kappa_{c,t}$ with a lower bound of 0.05.

$$T_{c,t} = T_{c,t-1} + \alpha \kappa_{c,t-1} (u_t - \gamma_1 T_{c,t-1}) + f m_t \quad (4.1)$$

$$m_t = m_{t-1} + \gamma_2 \sum_{c=\{A,B,\dots\}} \alpha (u_t - T_{c,t-1}) \quad (4.2)$$

Prediction errors for each context are computed similar to the standard TD-Momentum model. Associability for each context is updated at each time step as follows,

$$\kappa_{c,t+1} = (1 - \eta) \kappa_{c,t} + \eta |u_t - \gamma_1 T_{c,t-1}| \quad (4.3)$$

The associability weight parameter, $0 \leq \eta \leq 1$, controls how much historic prediction errors influence future associability values (this is general and not unique per context).

4.1.2 Implications of Associability TD-Momentum

As associability essentially creates unique, dynamic learning rates for each context, this model provides us with insights into how different contexts (traumatic/non-traumatic) influence learning dynamically over time. If threat learning were modulated by attention-based learning where traumatic contexts are more influential to learning, this model may explain why strong maladaptive cPTSD priors are so difficult to target via treatment.

Fig. 4.1 shows an example simulation with an input of 10 random attacks in context A (none in context B), associability for each context is initialised at 1 (similar to Brown et al. [7]). Associability of context A returns to the lower bound 0.05 rapidly after attacks, while associability of context B remains at the lower bound throughout. However, an example of the sensitised learning influence of traumatic context A is seen




Figure 4.1: Identical parameters and attack sequence to TD-Momentum Simulation in 4.2. Associability TD-Momentum threat (left) with associability weight parameter $\eta = 0.2$ and related associability values (right). Associability model generates smoother threat predictions than TD-Momentum, with threat being much similar between both contexts.

in the highlighted box. Two attacks occur here in quick succession, threat in context A increases more due to the second attack than threat in context B (left plot), this can be explained by associability. Context A associability κ_A does not reach 0 before the second attack occurs (right plot), causing a larger spike in associability for context A when the second attack is experienced. This results in a larger dynamic learning rate at this time, meaning that context A threat increases by a larger amount compared to context B threat.

The associability of each context can *indirectly* effect learning in all other contexts through the momentum term (computed from the decayed sum of prediction errors across *all* contexts). As traumatic, unpredictable contexts have larger associability values (thus larger dynamic learning rates), the prediction errors generated in these contexts will be more influential to learning compared to those generated in less traumatic, predictable contexts. Meaning they will contribute more to momentum, i.e. trauma in unpredictable contexts will influence momentum more than the same trauma in predictable contexts.

This may cause threat estimations in more predictable contexts to be heavily influenced (via momentum) by unpredictable contexts, which could explain why threat predictions for both contexts in our simulation are so similar (Fig. 4.1). Prolonged childhood trauma would strengthen the large overpowered associability values for trau-

matic contexts, as this would ensure they maintain high associability values, amplifying this effect, perhaps even bringing threat predictions in all contexts to the same level as the original trauma context.

This model provides a method for investigating the increased effect of salient traumatic contexts, *why* individuals with PTSD can be so affected by reminders of the context of the trauma, and may explain extinction learning failures in treatment. For cPTSD, it offers a potential explanation for issues with affect regulation (e.g. heightened emotional reactivity to small stressors/contextuals unrelated to original trauma). Associability creates a dynamic learning rate for each context that is the same for both positive and negative prediction errors within that context. Meaning, there is no difference between learning from positive or negative prediction errors in a given context. As learning in PTSD individuals may be modulated by unexpected experiences [14], incorporating separate learning rates for positive and negative prediction errors may lead to an improved representation of cPTSD learning. This is explored in the next section.

4.2 Incorporating Outcome-Sensitivity in TD-Momentum Model

Here we incorporate the “risk-sensitive” element of the RL model proposed by Ross et al. (section 2.1.3), although here we term the approach as “outcome-sensitive”. The motive here being that PTSD individuals may learn differently based on whether an outcome is better or worse than their prediction (i.e. the sign of the prediction error signifying good news or bad news). As such, splitting the single learning rate of the TD-Momentum model into separate positive and negative learning rates may provide an improved representation of learning. This model explores how threat learning may be modulated by how one perceives outcomes. Including momentum here incorporates non-associative learning, allowing for prediction errors in any context to influence threat in all other contexts.

4.2.1 Outcome-Sensitive TD-Momentum Model Form

The form is identical to the original model, with the learning rate used depending on the sign of the prediction error ($u_t - \gamma_1 T_{c,t-1}$). Separate learning rates are reflected within the momentum term where we have inserted the γ_1 decay rate to ensure prediction errors

are consistent between momentum calculation and the main update rule.

$$T_{c,t} = T_{c,t-1} + \alpha^{+/-} (u_t - \gamma_1 T_{c,t-1}) + f m_t \quad (4.4)$$

$$m_t = m_{t-1} + \gamma_2 \sum_{c=\{A,B,\dots\}} \alpha^{+/-} (u_t - \gamma_1 T_{c,t-1}) \quad (4.5)$$

Learning rate is: α^+ when the prediction error is positive (larger than expected outcome), and α^- when the prediction error is negative (smaller than expected outcome). Thus, large, unexpected threat outcomes cause positive prediction errors here, so α^+ relates to learning from very traumatising experiences.

4.2.2 Implications of Outcome-Sensitive TD-Momentum

Hauser et al. [13] implemented a similar outcome-sensitive RL model investigating developmental aspects of cognitive flexibility (the ability to adapt to unplanned events) in adolescents, they compared adolescents and adults who performed a probabilistic reversal learning task. They found that adolescents had increased sensitivity to *negative* prediction errors compared to adults (where the outcome is less than expected, i.e. a punishment). These findings may also translate to the threat learning context. Although, we expect that in this scenario, where we have assigned a positive learning rate to larger than expected outcomes (large threat or trauma), that learning is more sensitive to *positive* prediction errors, as these represent more intense, unexpected trauma. Fig. 4.2 shows a simulation of this model. The larger α^+ is used following underestimated outcomes (attacks) and causes large increases in threat estimation in comparison to the original TD-Momentum model.

By introducing positive and negative learning rates, momentum updates will vary based on whether the outcome prediction is over-estimated or under-estimated, leading to differences in how non-associative threat is learned compared to the standard TD-momentum model. Fig. 4.2 shows how threat in context B is larger than in the standard simulation, in line with the larger α^+ used. Thus, threat transferred through non-associative learning is influenced by these outcome-sensitive learning rates, and could lead to drastic variations of threat in unrelated contexts. This may have links to the increased reactivity to minor stressors unrelated to original trauma, typical of PTSD.

Investigating differences in learning sensitivity (positive and negative learning rates) between PTSD/control individuals who experience various threatening outcomes may provide insight on how learning mechanisms behave when outcomes are perceived

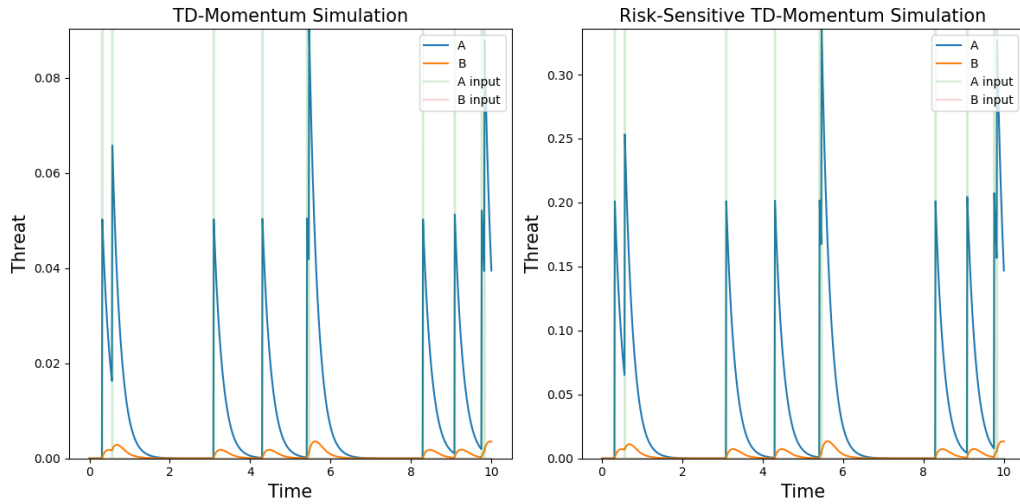


Figure 4.2: TD-Momentum plot (left) with parameters: $\alpha = 0.05, \gamma_1 = 0.9999, \gamma_2 = 0.05, f = 0.1$. Risk-Sensitive TD Momentum threat (right): $\alpha^- = 0.05, \alpha^+ = 0.2$. 10 random attacks across life in context A, none in B. Positive prediction errors use the larger α^+ causing larger overall threat and increases across both contexts (right).

differently by different groups. By reviewing any distinct group differences, this could provide the basis for a tool to determine PTSD risk and/or disorder trajectory.

A potential pitfall for this model (and the previous associability model) is the binary representation of inputs (0 for no attack, 1 for attack). This basic manipulation of threat sequences may result in lack of precision for fitted parameters. A model that can take a more continuous representation of threat as input allows for more complex threat sequences to be explored (e.g. between -1 and 1), and may provide more insight on fitting to behavioural data. A potential is explored in the final extension.

4.3 Incorporating Valence-Partitioning in TD-Momentum Model

Liebenow et al. [18] and Sands et al. [24] propose Valence Partitioning (VP) as a methodology for decoupling the algorithmic representation of rewards and punishments. They found that VP RL was effective at predicting dynamic changes in human choice behaviour and subjective experience. Consequently, they suggest that VP RL can be effective in deriving insights on mechanisms related to psychiatric disorders.

Standard TD learning treats punishment and reward as opposite ends of a single reward spectrum. Liebenow et al. propose this may not accurately reflect the true

processes of how learning (and associated behaviour) operates. VP TD learning maintains the successful process of TD learning, but implements two independent, parallel positive and negative valence systems. Liebenow et al. [18] apply a continuous outcome scale between -1 and +1, where -1 represents maximum punishment and +1 represents maximum reward. A valence threshold is set at 0, i.e. rewards are positively valenced and punishments are negatively valenced. Positively valenced outcomes are processed via the positively valenced system, and negatively valenced outcomes are processed via the negatively valenced system. These parallel systems allow for asymmetric representations of learning from positively and negatively valenced outcomes. We use VP RL here to capture any asymmetries in learning due to PTSD, exploring if threat learning is modulated by *how* individuals valence different outcomes.

4.3.1 Valence-Partitioned TD-Momentum Model Form

Determining which valence system processes the outcome, how prediction errors are computed, and which learning rate is used, depends only on the valence of the outcome. Valenced prediction errors are generated as shown in Eqs. (4.6) and (4.7). Outcome values are within $-1 \leq u_t \leq 1$, where 1 represents the most safe, pleasant outcomes, and -1 represents the most threatening outcomes (0 represents a null outcome). If outcome valence is not within the receptive field of a system (where we have set a *threshold* = 0), the outcome is treated as a null outcome for that system. Prediction error in context c at time t for each valence system is computed as follows:

$$\delta_{c,t}^P = \begin{cases} u_t + \gamma^P V_{c,t+1}^P - V_{c,t}^P & \text{if } u_t > 0 \\ 0 + \gamma^P V_{c,t+1}^P - V_{c,t}^P & \text{if } u_t \leq 0 \end{cases} \quad (4.6)$$

$$\delta_{c,t}^N = \begin{cases} |u_t| + \gamma^N V_{c,t+1}^N - V_{c,t}^N & \text{if } u_t < 0 \\ 0 + \gamma^N V_{c,t+1}^N - V_{c,t}^N & \text{if } u_t \geq 0 \end{cases} \quad (4.7)$$

Two value functions are generated, one for appetitive values (Positive system - Eq. (4.8) processes pleasant outcomes) and one for aversive values (Negative system - Eq. (4.9) processes threatening outcomes).

$$V_{c,t+1}^P = V_{c,t}^P + \alpha^P \cdot \delta_{c,t}^P + fm_t^P \quad (4.8)$$

$$V_{c,t+1}^N = V_{c,t}^N + \alpha^N \cdot \delta_{c,t}^N + fm_t^N \quad (4.9)$$

Separate momentum terms are created for positively and negatively valenced outcomes across all contexts. Update form remains the same, only we include the corresponding positive- or negative-specific prediction errors ($\delta_{c,t}^P$ or $\delta_{c,t}^N$) and learning rates (α^P or α^N) based on valence of the outcome.

$$m_t^P = m_{t-1}^P + \gamma_2 \sum_{c=\{A,B,\dots\}} \alpha^P \delta_{c,t}^P \quad (4.10)$$

$$m_t^N = m_{t-1}^N + \gamma_2 \sum_{c=\{A,B,\dots\}} \alpha^N \delta_{c,t}^N \quad (4.11)$$

Appetitive and aversive value functions are then combined to create the overall threat prediction in any context at any given time.

$$T_{c,t} = V_{c,t}^P - V_{c,t}^N \quad (4.12)$$

4.3.2 Implications of Valence-Partitioned TD-Momentum

Tracking the distinct valenced value functions and momentum terms will show us how behaviour of pleasantness/threat predictions across contexts varies between stimuli classed as “pleasant“ and “threatening“, i.e. positive and negative valence. Differences in group behaviour (PTSD/control) observed could provide explanations for how different stimuli can affect predictions over a lifetime and why symptom timescale and severity may differ between individuals (or groups). Regarding the distinct positively- and negatively-valenced momentums, we would expect the negatively valenced momentum to be larger and more influential to learning across all contexts as it represents the non-associative threat across all contexts from threatening outcomes. Comparing valenced momentum terms to the original single momentum will show, specifically, how much threatening outcomes can influence non-associative learning across all contexts, compared to positive outcomes.

A simulation of the VP TD-Momentum model with a larger negative learning rate is shown in Fig. 4.3, exposed to 10 positive (0.5, 1) and 10 negative outcomes (-1, -0.5), representing a broader spectrum of pleasant/threatening outcomes compared to previous extensions. This manipulation of outcomes leads to more diverse pleasant/threatening estimations in both contexts. Context B predictions (due to momentum) are small in value, oscillating around the neutral outcome of 0 due to the combination of positive and negative outcomes. The larger α^N parameter shows a disproportionate effect on learning from negatively valenced outcomes compared to positively valenced, shown by the downward spikes being more pronounced. The related VP value functions

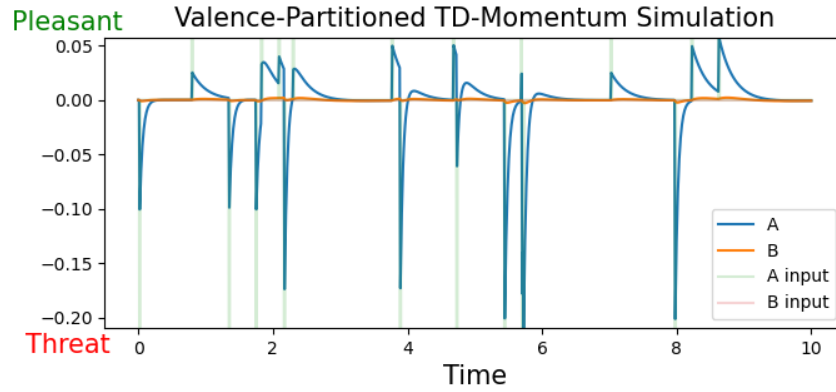


Figure 4.3: Valence-Partitioned TD-Momentum fit with identical parameters to TD Momentum Simulation in Fig. 4.2, and $\alpha^P = 0.05$, $\alpha^N = 0.2$. 10 positive (pleasant) outcomes and 10 negative (threat) stimuli. Larger negative learning rate α^- causes negative outcomes to influence learning more.

for both contexts, as well as the valenced momentum terms can be found in Figs. A.14, and A.15 of Appendix A. As expected, the value functions for context B are much lower in absolute value compared to context A; the negatively valenced value functions for both contexts are larger than the positively valenced, with the negatively valenced value function for trauma context A being disproportionately large (due to events occurring here) (Fig. A.14 left, orange). The negatively valenced momentum from threatening outcomes is larger and more influential to learning than the positively valenced momentum (Fig. A.15, orange) as expected and noted above. This is reflected in the overall fit on Fig. 4.3, negative valenced outcomes effect momentum more and thus are more influential to context B threat compared to positively valenced outcomes.

The *threshold* value of 0 may be subjective and investigations into fitting this to individuals/groups may provide insight on the heightened (or lessened, i.e. habituation) sensitivity to threats in PTSD. For example, this threshold value may be dysfunctionally large (e.g. 0.3) for highly symptomatic individuals, meaning they perceive small, pleasant outcomes as negative outcomes, processing them through the negatively valenced system, causing the negatively valenced system to influence learning more.

Although not applied by Liebenow et al. or Sands et al., we could include sensitivity parameters for positively and negatively valenced outcomes (β^P, β^N), similar to parameters proposed by Yanamori et al. (section 2.1.4). Overall pleasantness/threat prediction is then $T_{c,t} = \beta^P V_{c,t}^P - \beta^N V_{c,t}^N$. An index similar to the “reward-punishment sensitivity index” of Yanamori et al. could then be created for each individual, which we

label the “valence sensitivity index” β^P/β^N , giving a unique value to each participant corresponding to where they lie on the positive vs negative valence sensitivity spectrum. Higher values correspond to individuals whose learning is more influenced by positively valenced outcomes, while lower values correspond to those more influenced by negatively valenced outcomes. Further investigations into this VP TD-Momentum model may provide more complex insights into the underlying mechanisms of cPTSD.

4.4 TD-Momentum Extensions Summary

We have presented three different approaches to extending the TD-Momentum model, incorporating concepts from various RL models of associability, outcome-sensitivity, and valence partitioning in order to offer further insight into the disorder. The momentum term is maintained in all extensions, allowing for non-associative threat learning.

Associability (Brown et al. [7] and Homan et al. [14]) was integrated with the motive that trial-by-trial associability values, which are influenced by a context's history of prediction errors, may modulate learning in each context. Considering the heightened learning resulting from high associability cues found by Brown et al., (also shown in our simulation) this approach suggests that similar attention-based modulation of learning may exist in threat learning across various contexts.

The outcome-sensitive (Ross et al. [23]) integration explores how learning rates may vary between individuals for positive and negative prediction errors, where PTSD/cPTSD may be linked with heightened sensitivity to one set of prediction errors.

The integration of valence-partitioning (Liebenow et al. [18] and Sands et al. [24]) creates separately valenced value functions positively- and negatively-valenced outcomes which are combined to create an overall pleasantness/threat prediction. This aims to capture how asymmetries in learning may be based on how individuals valence outcomes differently.

By applying results of relevant behavioural tasks performed by control and disorder-affected participants we can analyse variations of parameters and model behaviour between groups. Applying such behavioural data to these proposed extensions is thus key in indicating how much value they can add with regards to explaining the disorder, and will reveal how associability, outcome-sensitivity, and valence partitioning can influence threat prediction learning.

Chapter 5

Conclusion

The aim of this dissertation has been to re-implement the TD-Momentum model proposed by Kaye et al. [16], assess the capability of this model for explaining mechanisms of cPTSD, and provide potential model extensions that can add value and further insight.

We re-implemented the Bayesian baseline model proposed by Kaye et al. and found it predicted the probability of attack well, showing how a simulated ideal observer may perform in a Bayesian setting when predicting attack rate over a lifetime. Through ELS scenarios we showed how various types of trauma in early life can cause a disproportionate affect on threat prediction, with clustered attacks in early life causing larger mean estimates and variance in predictions.

TD and TD-Momentum RL models were then re-implemented and fit to data collected by an SEFL experiment. Mouse freezing data was fit and the model comparison showed that stressed mice favoured the TD-Momentum model while unstressed mice favoured the simpler TD model. We concluded similar to Kaye et al., that the stressed mice prefer the inclusion of the momentum term due to the incorporation of non-associative learning, i.e. momentum allows for threat predictions from the stressful day 1 context A to influence predictions across days 6 and 7 in context B.

Following this, we extended the findings of Kaye et al. by performing our own investigations into TD-Momentum model behaviour and how ELS scenarios affect predictions across contexts. We suggest that the oscillatory behaviour built into the momentum term due to certain parameter values may reflect differences between healthy and highly symptomatic individuals, where more oscillatory behaviour in momentum may have links to issues with affect regulation typically seen in cPTSD. The ELS scenarios showed that clustered attacks experienced in early life seem to have a larger disproportionate impact on threat learning within the original trauma context (similar

to the Bayesian model), as well as in novel contexts via momentum. We found that clustered attacks administered in early life may also result in habituation to future attacks, which may be linked with the emotional numbing effects seen in cPTSD.

We also found similar results to Kaye et al. in relation to how extinction learning can be negatively affected by the momentum term, and may affect treatments like exposure therapy by generating a sensitised effect, rather than habituation. We suggest that a maintained high sense of threat in a recently exited trauma context, due to increased momentum, may cause threat in other contexts to be maintained at a higher rate than usual, resulting in symptom clusters such as increased reactivity/arousal and avoidance which may lead to difficulties in various aspects of life.

After exploring how much the TD-Momentum model can explain we offered three potential extensions which may add value and provide us with a tool for investigating symptoms, individual risk and implications to treatments. Proposed methodologies included: an associability term to act as a context-dependent momentum, the inclusion of outcome-sensitivity to assign different learning rates for positive/negative prediction errors, and the idea of partitioning positive (pleasant outcomes) and negative (threatening outcomes) learning based on valence of the outcome at each timestep.

These extensions aim to combat some of the pitfalls of the standard TD-Momentum model, such as: the heightened influence to learning that some (traumatic) contexts may have (associability model addresses this), the uniform view of how different outcomes may influence prediction more/less (constant learning rate for all outcomes, outcome-sensitive model addresses this), and the fact that learning from safe, pleasant outcomes and threatening outcomes may be processed by entirely different systems, based on how individuals interpret outcomes (valence partitioning addresses this).

Ultimately, further investigation into these extensions (and similar computational models) will improve the understanding of the underlying mechanisms of cPTSD, which is vital in defining the distinct differences and similarities between PTSD and other psychiatric disorders. These findings will be key in discovering the best methods for risk assessment, treatment planning and mapping trajectory of the disorder. Although computational psychiatry is still at a youthful stage, there is an increasing amount of research being published in this field. The findings from such research hope to someday (in the near future) be applied for clinical uses to improve the quality of life for all those affected by such mental health problems.

Bibliography

- [1] Python Package Index - PyPI. Publisher: Python Software Foundation.
- [2] Robin Achterhof, Rafaële J. C. Huntjens, Marie-Louise Meewisse, and Henk A. L. Kiers. Assessing the application of latent class and latent profile analysis for evaluating the construct validity of complex posttraumatic stress disorder: cautions and limitations. *European Journal of Psychotraumatology*, 10(1):1698223, December 2019.
- [3] Laurence Aitchison and Máté Lengyel. With or without you: predictive coding and Bayesian inference in the brain. *Current Opinion in Neurobiology*, 46:219–227, October 2017.
- [4] American Psychiatric Association. *Diagnostic and statistical manual of mental disorders (5th ed.)*. American Psychiatric Publishing, 2013.
- [5] Jonathan I Bisson. Post-traumatic stress disorder. *BMJ*, 334(7597):789–793, April 2007.
- [6] Kathleen T Brady, Tim Brewerton, Therese Killeen, and Sylvia Lucerini. Comorbidity of Psychiatric Disorders and Posttraumatic Stress Disorder. *J Clin Psychiatry*, April 2000.
- [7] Vanessa M Brown, Lusha Zhu, John M Wang, B Christopher Frueh, Brooks King-Casas, and Pearl H Chiu. Associability-modulated loss learning is increased in posttraumatic stress disorder. *eLife*, 7:e30150, January 2018.
- [8] Marylène Cloitre, Chris R. Brewin, Jonathan I. Bisson, Philip Hyland, Thanos Karatzias, Brigitte Lueger-Schuster, Andreas Maercker, Neil P. Roberts, and Mark Shevlin. Evidence for the coherence and integrity of the complex PTSD (CPTSD) diagnosis: response to Achterhof et al., (2019) and Ford (2020). *European Journal of Psychotraumatology*, 11(1):1739873, December 2020.

- [9] Eran Eldar and Yael Niv. Interaction between emotional state and learning underlies mood instability. *Nature Communications*, 6(1):6149, January 2015.
- [10] Eran Eldar, Robb B. Rutledge, Raymond J. Dolan, and Yael Niv. Mood as Representation of Momentum. *Trends in Cognitive Sciences*, 20(1):15–24, January 2016.
- [11] Daniel Foreman-Mackey, David W. Hogg, Dustin Lang, and Jonathan Goodman. emcee : The MCMC Hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306–312, March 2013.
- [12] Jonathan Goodman and Jonathan Weare. Ensemble samplers with affine invariance. *Communications in Applied Mathematics and Computational Science*, 5(1):65–80, January 2010.
- [13] Tobias U. Hauser, Reto Iannaccone, Susanne Walitza, Daniel Brandeis, and Silvia Brem. Cognitive flexibility in adolescence: Neural and behavioral mechanisms of reward prediction error processing in adaptive decision making during development. *NeuroImage*, 104:347–354, January 2015.
- [14] Philipp Homan, Ifat Levy, Eric Feltham, Charles Gordon, Jingchu Hu, Jian Li, Robert H. Pietrzak, Steven Southwick, John H. Krystal, Ilan Harpaz-Rotem, and Daniela Schiller. Neural computations of threat in the aftermath of combat trauma. *Nature Neuroscience*, 22(3):470–476, March 2019.
- [15] Eric Jones, Travis Oliphant, Pearu Peterson, and others. SciPy: Open source scientific tools for Python, 2001.
- [16] Alfred P. Kaye, Manasa G. Rao, Alex C. Kwan, Kerry J. Ressler, and John H. Krystal. A computational model for learning from repeated traumatic experiences under uncertainty. *Cognitive, Affective, & Behavioral Neuroscience*, May 2023.
- [17] Ye Ji Kim, Sanne J.H. Rooij, Timothy D. Ely, Negar Fani, Kerry J. Ressler, Tanja Jovanovic, and Jennifer S. Stevens. Association between posttraumatic stress disorder severity and amygdala habituation to fearful stimuli. *Depression and Anxiety*, 36(7):647–658, July 2019.
- [18] Brittany Liebenow, Rachel Jones, Emily DiMarco, Jonathan D. Trattner, Joseph Humphries, L. Paul Sands, Kasey P. Spry, Christina K. Johnson, Evelyn B. Farkas,

- Angela Jiang, and Kenneth T. Kishida. Computational reinforcement learning, reward (and punishment), and dopamine in psychiatric disorders. *Frontiers in Psychiatry*, 13:886297, October 2022.
- [19] Milen L. Radell, Catherine E. Myers, Jony Sheynin, and Ahmed A. Moustafa. Computational Models of Post-traumatic Stress Disorder (PTSD). In *Computational Models of Brain and Behavior*, pages 43–55. John Wiley & Sons, Ltd, 2017. Section: 4 _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119159193.ch4>.
- [20] Vinuta Rau, Joseph P. DeCola, and Michael S. Fanselow. Stress-induced enhancement of fear learning: An animal model of posttraumatic stress disorder. *Neuroscience & Biobehavioral Reviews*, 29(8):1207–1223, 2005.
- [21] Robert Rescorla. *A theory of Pavlovian conditioning : Variations in the effectiveness of reinforcement and nonreinforcement*. Appleton-Century-Crofts, New York, 1972.
- [22] Kerry. J. Ressler, Sabina Berretta, Vadim Y. Bolshakov, Isabelle M. Rosso, Edward G. Meloni, Scott L. Rauch, and William A. Carlezon. Post-traumatic stress disorder: clinical and translational neuroscience from cells to circuits. *Nature Reviews Neurology*, 18(5):273–288, May 2022.
- [23] Marisa C. Ross, Jennifer K. Lenow, Clinton D. Kilts, and Josh M. Cisler. Altered neural encoding of prediction errors in assault-related posttraumatic stress disorder. *Journal of Psychiatric Research*, 103:83–90, August 2018.
- [24] L. Paul Sands, Angela Jiang, Rachel E. Jones, Jonathan D. Trattner, and Kenneth T. Kishida. Valence-partitioned learning signals drive choice behavior and phenomenal subjective experience in humans. preprint, Neuroscience, March 2023.
- [25] Skipper Seabold and Josef Perktold. statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*, 2010.
- [26] Richard S. Sutton and Andrew G. Barto. *Reinforcement learning: an introduction*. Adaptive computation and machine learning series. The MIT Press, Cambridge, Massachusetts, second edition edition, 2018.

- [27] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.
- [28] Christopher Watkins and Peter Dayan. Q-Learning. *Kluwer Academic Publishers*, pages 279–292, 1992.
- [29] World Health Organization (WHO). *International Classification of Diseases, Eleventh Revision (ICD-11)*. (WHO), 2019.
- [30] Sam Wilkinson, Guy Dodgson, and Kevin Meares. Predictive Processing and the Varieties of Psychological Trauma. *Frontiers in Psychology*, 8:1840, October 2017.
- [31] Yumeya Yamamori, Oliver J Robinson, and Jonathan P Roiser. Approach-avoidance reinforcement learning as a translational and computational model of anxiety-related avoidance. preprint, *elife*, June 2023.

Appendix A

Additional Plots

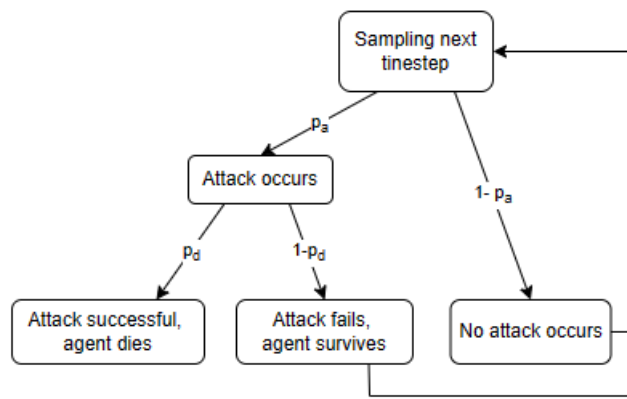


Figure A.1: Bayesian Probabilistic model of trauma attack estimation (re-created from Kaye et al. [16]). Attacks continue until agent death or end of timesteps.




Figure A.2: Variance in p_a estimation decreases over the course of a lifetime, showing how the Bayesian model progressively improves and becomes more confident in its estimates. This plot shows the variance in p_a over a lifetime for a typical sequence, similar to that used in Fig. 3.1 (related code in Appendix B.4)

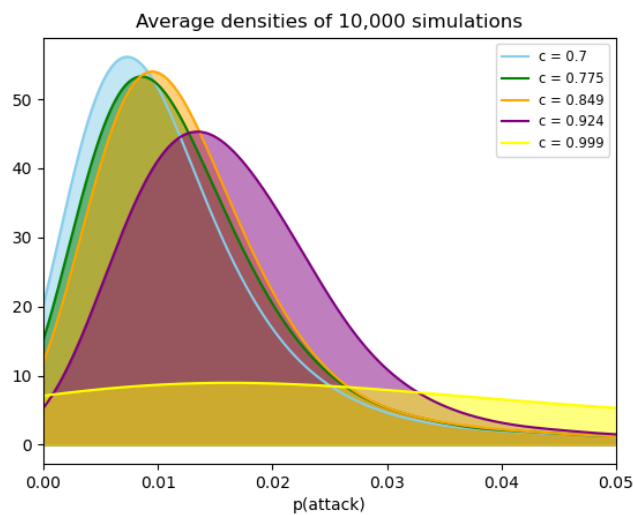


Figure A.3: Densities “averaged” over 10,000 simulations of AR time series for each correlation value. Larger AR correlation results in larger variance for threat predictions of p_a . Sharper, more precise predictions are given by $c = 0.7$, while flatter, less precise predictions given by $c = 0.999$.

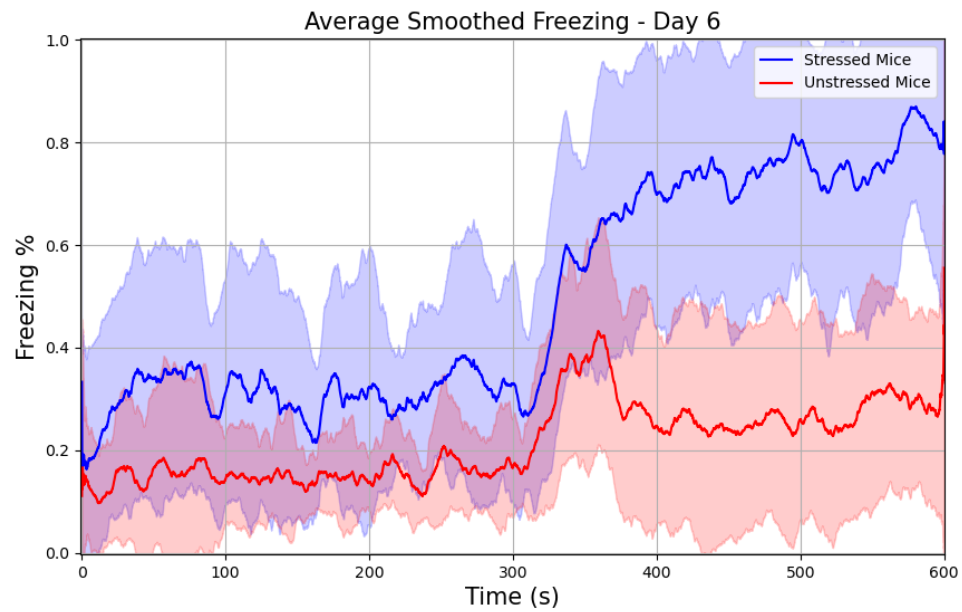


Figure A.4: Average freezing of 15 stressed and 18 unstressed mice across day 6 of the SEFL experiment.




Figure A.5: Average freezing of 15 stressed and 18 unstressed mice across day 7 of the SEFL experiment.

TD Model - Multivariate Sampled Parameters - Recovery

α Coeff	α P-Val	γ_1 Coeff	γ_1 P-Val
0.9994	0.0000	0.8775	0.0000
0.9993	0.0000	0.9647	0.0000
1.0000	0.0000	0.9993	0.0000
1.0000	0.0000	0.9299	0.0000
0.9986	0.0000	0.9345	0.0000

Figure A.6: TD model parameter recovery Pearson correlation coefficients and p-values for multivariate sampled parameter values

TD Model - List Sampled Parameters - Recovery

α Coeff	α P-Val	γ_1 Coeff	γ_1 P-Val
1.0000	0.0000	0.9999	0.0000
1.0000	0.0000	1.0000	0.0000
1.0000	0.0000	1.0000	0.0000
1.0000	0.0000	0.9999	0.0000
1.0000	0.0000	0.9999	0.0000

Figure A.7: TD model parameter recovery Pearson correlation coefficients and p-values for list sampled parameter values

TD-Momentum Model - Multivariate Sampled Parameters - Recovery

α Coeff	α P-Val	γ_1 Coeff	γ_1 P-Val	γ_2 Coeff	γ_2 P-Val	f Coeff	f P-Val
0.9952	0.0000	0.6501	0.0000	0.4513	0.0010	0.6032	0.0000
0.9904	0.0000	0.8102	0.0000	0.5473	0.0000	0.4103	0.0031
0.9968	0.0000	0.5721	0.0000	0.4032	0.0037	0.3517	0.0123
0.9906	0.0000	0.8645	0.0000	0.4673	0.0006	0.2117	0.1400
0.9951	0.0000	0.7844	0.0000	0.4486	0.0011	0.3493	0.0129

Figure A.8: TD-Momentum model parameter recovery Pearson correlation coefficients and p-values for multivariate sampled parameter values

TD-Momentum Model - List Sampled Parameters - Recovery

α Coeff	α P-Val	γ_1 Coeff	γ_1 P-Val	γ_2 Coeff	γ_2 P-Val	f Coeff	f P-Val
0.9986	0.0000	0.8007	0.0000	0.3877	0.0054	0.4329	0.0017
0.9992	0.0000	0.8618	0.0000	0.5386	0.0001	0.2829	0.0465
0.9983	0.0000	0.6341	0.0000	0.6814	0.0000	0.5150	0.0001
0.9963	0.0000	0.6160	0.0000	0.6156	0.0000	0.4618	0.0007
0.9979	0.0000	0.6241	0.0000	0.4856	0.0004	0.4397	0.0014

Figure A.9: TD-Momentum model parameter recovery Pearson correlation coefficients and p-values for list sampled parameter values

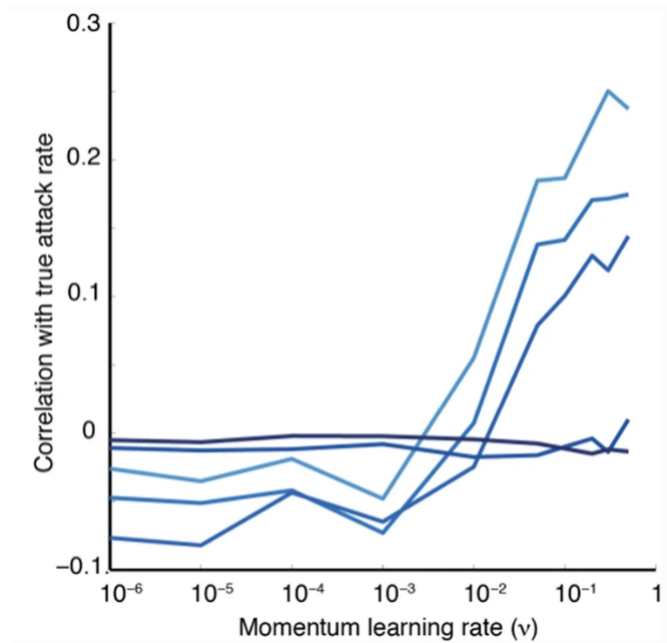


Figure A.10: Variations of auto-correlation in the AR process producing attack rates used to generate attack sequences shows how including momentum assists in extracting more information about the true attack rate (light blue indicates highest autoregression, dark blue indicates lowest autoregression) (figure by Kaye et al. [16])




Figure A.11: Associability TD-Momentum model with $\eta = 0.001$. Very low values of η reduce this extended model to the original TD-Momentum model.

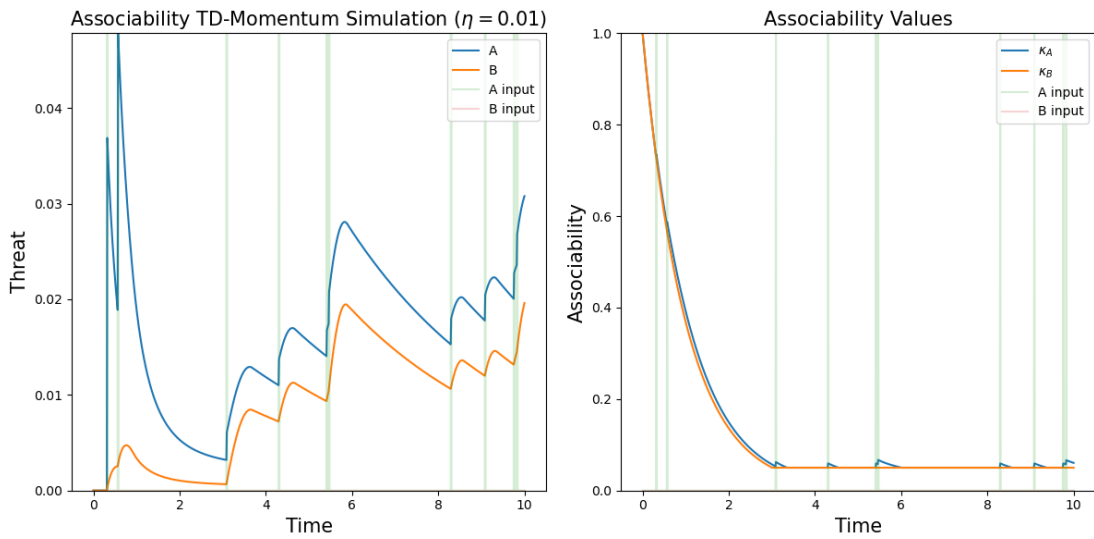


Figure A.12: Associability TD-Momentum model with $\eta = 0.01$. Increasing the value of η changes the shape of threat predictions across both contexts, reducing overall threat predictions and creating smoother predictions




Figure A.13: Associability TD-Momentum model with $\eta = 0.9999$. Very high values of η generate threat predictions for both contexts that are very similar. Context B threat is only slightly below that of the trauma context A.

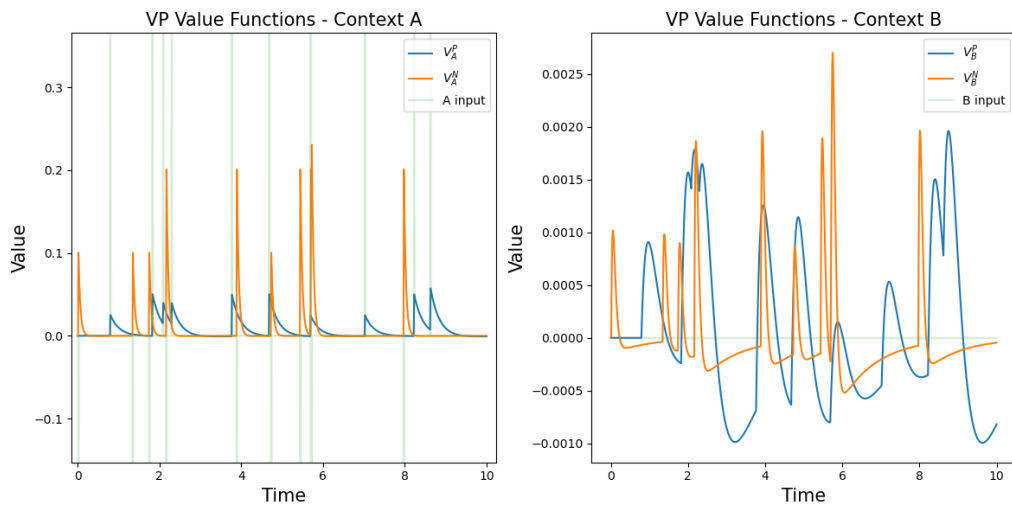


Figure A.14: Individual positive and negative VP value functions for each context (left - A, right - B). These are Eq. (4.8) and Eq. (4.9).

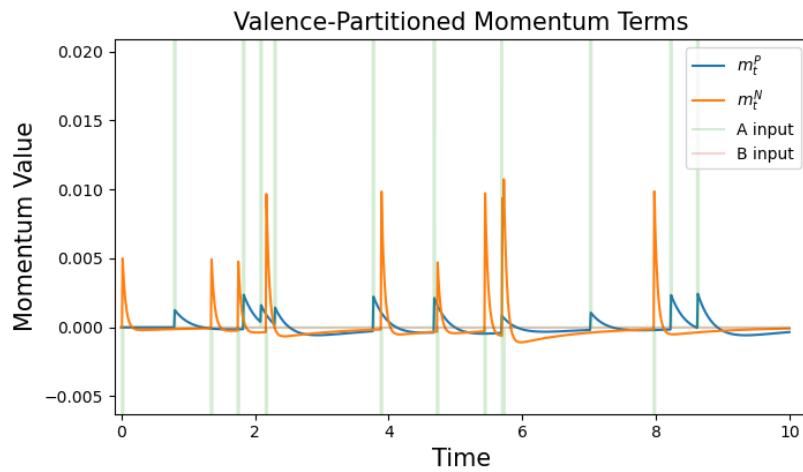


Figure A.15: Individual positively and negatively valenced momentum functions (blue - positive, orange - negative) relating to Eq. (4.10) and Eq. (4.11). Note how the negatively valenced momentum, containing the prediction errors of negative trauma outcomes, is larger and more influential than the positively valenced momentum of positive, non-traumatic outcomes.

Appendix B

Bayesian Model Python Code

B.1 Bayesian Model - Attack & Death Simulations

```
1 import numpy as np
2 import pylab as pl
3 import scipy as sp
4 import scipy.stats as stats
5 import matplotlib.pyplot as plt
6 import matplotlib; matplotlib.use('Qt5Agg')
7 import emcee
8 from emcee import moves
9 import random
10 from scipy.ndimage import gaussian_filter
11 from matplotlib.colors import Normalize
12 from statsmodels.tsa.arima_process import ArmaProcess
13
14 def sim_sequences(pa, pd, num_time_steps):
15     '''
16     Simulates attack and death sequences based on binomial
17     distribution, with constant pa and pd across timesteps.
18
19     :param prob_attack: Probability of attack
20     :param prob_death: Probability of death given attack happens
21     :param num_time_steps: Total number of trials or timesteps where
22     attack can occur
23
24     :return: attacks, deaths which are arrays corresponding to
25     sequences of attacks and deaths (0 is no attack/no death,
26     1 is attack/death)
```

```

24     '''
25     # Simulating attack sequence where attacks are binomially
    distributed
26     attacks = np.random.binomial(1, pa, num_time_steps)
27
28     # Death sequence based on probab of death given attack
29     deaths = np.random.binomial(1, pd, num_time_steps) * attacks
30
31     return attacks, deaths

```

B.2 Bayesian Model - MCMC Sampler

```

1  ### Bayesian model fitting procedure ###
2  # Now we move onto the Bayesian model fitting procedure.
3
4  # 'gwmcmc' function in matlab, here we use the emcee package in
    Python -
5  # Goodman and Weares Affine Invariant Markov chain Monte Carlo (MCMC
    ) Ensemble sampler
6
7  def ln_likelihood(params, attacks, deaths):
8      '''
9      Computes natural log of the joint likelihood
10
11      :param params: Values for pa, pd
12      :param attacks: Attack sequence
13      :param deaths: Death sequence
14
15      :return: Log-likelihood
16      '''
17      pa, pd = params
18
19      # Compute number of attacks and deaths
20      num_attacks = np.sum(attacks)
21      num_deaths = np.sum(attacks * deaths)
22
23      # Calculate log-likelihood
24      ln_likelihood = num_attacks * np.log(pa * (1 - pd)) + (len(
        attacks) - num_attacks) * np.log(1 - pa) + num_deaths * np.log(pa
        * pd)
25

```

```
26     return ln_likelihood
27
28 def ln_prior(params):
29     '''
30     Computes natural log of prior
31
32     :param params: Values for pa, pd
33
34     :return: Log-prior
35     '''
36     pa, pd = params
37
38     # Flat prior (uniform distribution), same for both pa and pd as
39     both are probabilities
40     p_min = 0 # Lower limit on range of values
41     p_max = 1 # Upper limit on range of values
42
43     # If pa, pd between correct range, 0-1, return prior = 1.0 (log-
44     prior = 0), else 0 (log-prior = -inf)
45     if p_min < pa < p_max and p_min < pd < p_max:
46         return 0.0
47     else:
48         return -np.inf
49
50 def ln_posterior(params, attacks, deaths):
51     '''
52     Computes natural log of the joint posterior
53
54     :param params: Values for pa, pd
55     :param attacks: Attack sequence
56     :param deaths: Death sequence
57
58     :return: Log-posterior
59     '''
60     # Compute prior from function above
61     ln_prior_val = ln_prior(params)
62
63     # If function is NOT finite, return prob of 0 (log-prior = -inf)
64     if not np.isfinite(ln_prior_val):
65         return -np.inf
```

```
66     # Posterior = Prior * Likelihood (Log-Prior + Log-Likelihood)
67     ln_posterior_val = ln_prior_val + ln_likelihood(params, attacks,
68     deaths)
69
70     return ln_posterior_val
71
72 def fit_bayesian_model(attacks, deaths, num_walkers, num_steps,
73     burn_in, step_size):
74     '''
75     Fits Bayesian model using Affine Invariant MCMC Sampler from
76     emcee package.
77
78     :param attacks: Attack sequence
79     :param deaths: Death sequence
80     :param num_walkers: Number of walkers
81     :param num_steps: Number of steps
82     :param burn_in: Percentage of initial steps to remove as burn-in
83     :param step_size: Step size to use in Stretch_move parameter of
84     sampler
85
86     :return: Samples from chain created by sampler, i.e. progressive
87     estimates for pa and pd
88     '''
89     # Define number of dims, i.e. pa and pd
90     num_dimensions = 2
91
92     # Initialise positions for the walkers
93     initial_positions = np.random.rand(num_walkers, num_dimensions)
94
95     # Set up the MCMC sampler with StretchMove - corresponds to "
96     stepsize" = 2 in paper (matlab equivalent)
97     stretch_move = moves.StretchMove(a=step_size)
98     sampler = emcee.EnsembleSampler(num_walkers, num_dimensions,
99     ln_posterior, args=(attacks, deaths),
100     moves=[stretch_move])
101
102     # Running burn-in phase, throwing these away so use _ as
103     placeholder variables
104     num_burn_in_steps = int(burn_in * num_steps)
105     _, _, _ = sampler.run_mcmc(initial_positions, num_burn_in_steps,
106     progress=True)
```

```

99
100     # Resetting the sampler
101     sampler.reset()
102
103     # Running production phase
104     sampler.run_mcmc(None, num_steps, progress=True)
105
106     # Retrieving MCMC samples from chain
107     samples = sampler.get_chain(discard=num_burn_in_steps, flat=True
108                                )
109
110     return samples

```

B.3 Bayesian Model - Posterior Plots

```

1 def bayesian_plots(samples):
2     '''
3     Takes output of fit_bayesian_model and recreates plots for fig 2
4     C in paper. Uses kernel density estimation (kde)
5     to smoothen posterior distributions for pa and pd.
6
7     :param samples: Output of fit_bayesian_model (estimations of pa
8     and pd from MCMC Sampler)
9
10    :return: Figure of 3 plots, p(attack), p(death|attack), p(attack
11    ) vs p(death|attack)
12    '''
13    # Create figure and axes
14    fig, axes = plt.subplots(2, 2, sharex='col')
15
16    # First plot - p(attack)
17    kdel = stats.gaussian_kde(samples[:, 0], bw_method=0.2)
18    x1 = np.linspace(0, 0.025, 1000)
19    y1 = kdel(x1)
20    axes[0, 0].fill_between(x1, y1, color='skyblue', alpha=0.5)
21    axes[0, 0].plot(x1, y1, label='p(attack) posterior', color='
22    skyblue')
23    axes[0, 0].set_xlim(0, 0.025)
24    axes[0, 0].set_ylabel('p(attack)')
25    axes[0, 0].legend(loc='upper right', fontsize='small')

```

```
23     # Second plot - contour plot of p(attack) against p(death|attack
24     )
25     hist, xedges, yedges = np.histogram2d(samples[:, 0], samples[:,
26     1], bins=25)
27     smooth_hist = gaussian_filter(hist.T, sigma=1)
28     axes[1, 0].contourf(xedges[:-1], yedges[:-1], smooth_hist, cmap=
29     'Blues')
30     #axes[1, 0].set_xlim(0, 0.025)
31     # Determine the minimum and maximum values for the x-axis
32     x_min = np.min(samples[:, 0])
33     #x_max = np.max(samples[:, 0])
34
35     # Set the x-axis limit based on the data range
36     axes[1, 0].set_xlim(x_min, 0.025)
37
38     axes[1, 0].set_ylabel('p(death|attack)')
39     axes[1, 0].set_xlabel('p(attack)')
40
41     # Third plot - p(death|attack)
42     kde3 = stats.gaussian_kde(samples[:, 1], bw_method=0.5)
43     x3 = np.linspace(0, 0.8, 1000) # changed from 1000 to 60
44     y3 = kde3(x3)
45     axes[1, 1].fill_between(x3, y3, color='skyblue', alpha=0.5)
46     axes[1, 1].plot(x3, y3, label='p(death|attack) posterior', color
47     ='skyblue')
48     axes[1, 1].set_xlabel('p(death|attack)')
49     axes[1, 1].legend(loc='upper right', fontsize='small')
50
51     # Remove top right empty plot
52     fig.delaxes(axes[0, 1])
53
54     # Adjust spacing between subplots
55     plt.tight_layout()
56     plt.show()
57
58 attacks, deaths = sim_sequences(0.01, 0.2, 700)
59 num_walkers = 30
60 num_steps = 60
61 burn_in = 0.3
62 step_size = 2
```



```

61 samples = fit_bayesian_model(attacks, deaths, num_walkers, num_steps
    , burn_in, step_size)
62 bayesian_plots(samples)

```

B.4 Bayesian Model - Variance in Threat Over Time

```

1 def variance_over_time(samples, colour, x_offset):
2     '''
3     Takes output of fit_bayesian_model, i.e. samples from MCMC
4     sampler, and returns plot showing the variance in
5     estimates over time. Groups use all available data up until the
6     end of the respective group.
7
8     :param samples: Output of fit_bayesian_model, samples from MCMC
9     sampler
10
11    :param colour: Colour scheme for plot
12
13    :param x_offset: Use when comparing more than 1 sequence,
14    separates plots for clarity
15
16    :return: Plot showing variance in estimates over lifetime
17    '''
18    x = np.linspace(0, 1, np.size(samples[:, 0]))
19    y = samples[:, 0]
20
21    # Divide sample pa's into groups to plot with error bars
22    num_groups = 8
23    group_size = len(y) // num_groups
24
25    # Empty sets to store group means and stds
26    means = []
27    stds = []
28
29    # Compute mean and std dev for each group - each group will use
30    info up until the last index of the group
31    for i in range(1, num_groups + 1):
32        end_idx = i * group_size
33        group = y[:end_idx]
34        means.append(np.mean(group))
35        stds.append(np.std(group))

```

```

31     # Compute the x values for the last of each group
32     group_last_x = x[group_size - 1::group_size]
33
34     # Add the x-axis offset - only used for plotting two sequences
35     # together, e.g. the ELS and lifetime scenarios
36     group_last_x += x_offset
37
38     # Compute colors based on x-values
39     cmap = plt.get_cmap(colour)
40     norm = Normalize(vmin=np.min(x), vmax=np.max(x))
41     colors = cmap(norm(group_last_x))
42
43     # Plot the error bars with colormap
44     for x_val, mean, std, color in zip(group_last_x, means, stds,
45     colors):
46         plt.errorbar(x_val, mean, yerr=std, fmt='o', color=color,
47         alpha=0.5, label='95% CI')
48
49     # Add colors to the points
50     sc = plt.scatter(group_last_x, means, c=colors, cmap=cmap)
51
52     # Add legend, labels, limits
53     plt.ylim(0, 0.6)
54     plt.xlabel('Time (fraction of lifetime)', fontsize=15)
55     plt.ylabel('Estimated attack rate', fontsize=15)
56     plt.title('Variance in  $p_a$  decreasing over time', fontsize=15)
57     plt.xticks(fontsize=14)
58     plt.yticks(fontsize=14)
59     plt.show()
60
61 samples = fit_bayesian_model(attacks, deaths, num_walkers, num_steps
62     , burn_in, step_size)
63
64 variance_over_time(samples, colour = 'cool', x_offset=0)

```

B.5 Bayesian Model - AR Time Series

```

1 def AR_PROCESS(c):
2     '''

```

```

3     Uses ArmaProcess function in tsa package to create a
4     autocorrelated attack rate time series from an AR process.
5
6     :param c: Chosen correlation coefficient to be used
7
8     :return: Attack rate time series of 700 time steps.
9     '''
10    ar_coeff = np.array([1, -c])    # coeffs for AR process
11    ma_coeff = np.array([1])       # equivalent to [1, 0], i.e. no
12    moving average, only AR process
13    AR_PROCESS = ArmaProcess(ar_coeff, ma_coeff).generate_sample(
14    nsample=700, scale=0.01)      # scale param is noise std
15
16    # Clip the AR process between 0 and 1
17    clipped_AR_PROCESS = np.clip(AR_PROCESS, 0, 1)
18
19    return clipped_AR_PROCESS

```

B.6 Bayesian Model - Posterior Comparisons for Various c Coefficients in AR Time Series

```

1 def average_density_many_ts(c, num_runs):
2     '''
3     Creates "average" density for many runs of time series. i.e.
4     creates an "average" of the plots created in the
5     bayesian_plots function over num_runs number of simulations of
6     time series.
7
8     :param c: Chosen correlation coefficient to be used
9     :param num_runs: Number of runs to compute "average" density
10    over
11
12    :return: Outputs stacked densities over amount of numruns, plots
13    "average" posterior in bayesian_plots
14    '''
15    # Create an empty array to store the results
16    results = np.empty((num_runs, 1260, 2))    # (1260, 2) is shape
17    of output from estimates_from_ts
18
19    # Run the function multiple times

```

```
15     for i in range(num_runs):
16         results[i] = estimates_from_ts(c)
17
18     # Calculate the average DENSITY of the results array
19     # Correct way to take mean over densities, not the variables:
20     stacked_values = np.column_stack((np.ravel(results[:, :, 0]), np
21     .ravel(results[:, :, 1])))
22     bayesian_plots(stacked_values)
23
24     return stacked_values
25
26 average = average_density_many_ts(c = 0.7, num_runs = 10)
27
28 # Here, we turn to the supplementary material word doc, and recreate
29 # supp fig 2 and 3A
30
31 # Supp Fig 2C: Dispersion of estimated pa over time for varying
32 # correlation coefficients (c=0.7 - 0.99)
33 # Additional - Supp Fig 2C - Dispersion of estimated attack rate for
34 # varying correlation (c) values:
35
36 n = 10000
37 c_values = np.linspace(0.7, 0.999, 5)
38
39 variances = np.empty_like(c_values)
40 std_devs = np.empty_like(c_values)
41 densities = {}
42
43 for t, c in enumerate(c_values):
44     dens = average_density_many_ts(c, n)[: ,0]
45     densities[t] = dens
46     variances[t] = np.std(dens)
47     std_devs[t] = np.var(dens)
48
49 # This provides us with a different view on sup fit 2C, here we've
50 # performed 10,000 sims per auto-correlation value
51 # We can see that the largest value c=0.999 does indeed provide us
52 # with the largest standard deviation and variance
53 # for estimated attack rate across all 10,000 runs. All other c vals
54 # are quite similar, indicating that c vals closer
55 # to 1 provide more variability in attack rate estimation.
```

```

50 # Average pa posteriors for 10,000 sims on each c value:
51
52 def posterior_comparison(c_vals, estimates):
53     '''
54     Takes output of average_density_many_ts function, i.e. 10,000
55     sims of various c values and creates pa posterior
56     plot for comparison. Shows increased variability in estimation
57     for larger c values.
58
59     :param c_vals: List of associated c values
60     :param estimates: Dictionary containing densities of each c
61     value generated via average_density_many_ts
62
63     :return: Plot comparing pa estimations
64     '''
65     colors = ['skyblue', 'green', 'orange', 'purple', 'yellow'] #
66     List of colors for each plot
67     labels = ['c = {}'.format(np.round(t, 3)) for t in c_vals] #
68     Labels for the legend
69
70     for t in range(len(estimates)):
71         # Posterior for p(attack)
72         kde1 = stats.gaussian_kde(estimates[t], bw_method=0.2)
73         x1 = np.linspace(0, 0.05, 1000)
74         y1 = kde1(x1)
75         plt.fill_between(x1, y1, color=colors[t], alpha=0.5)
76         plt.plot(x1, y1, label=labels[t], color=colors[t])
77
78     plt.xlim(0, 0.05)
79     plt.xlabel('p(attack)')
80     plt.title('Average densities of 10,000 simulations')
81     plt.legend(loc='upper right', fontsize='small')
82
83     plt.show()
84
85 posterior_comparison(c_values, densities)

```

B.7 Bayesian Model - ELS Example

```

1 def sim_scenarios(pd, num_time_steps, num_attacks):
2     '''

```

```

3     Simulates attack and death sequences for two scenarios,
4     random_lifetime and random_els. Limiting the number
5     of attacks to be equal for both scenarios = num_attacks.
6
7     :param probab_attack: Probability of attack
8     :param probab_death: Probability of death given attack happens
9     :param num_time_steps: Total number of trials or timesteps where
10    attack can occur
11
12    :return: attacks, deaths which are arrays corresponding to
13    sequences of attacks and deaths (0 is no attack/no death,
14    1 is attack/death)
15    '''
16
17    # Lifetime Random Scenario:
18    rand_life_attacks = np.zeros(num_time_steps)
19    life_indices = np.random.choice(num_time_steps, num_attacks,
20    replace=False)
21
22    # Input attacks (1s) in random chosen timesteps
23    rand_life_attacks[life_indices] = 1
24    rand_life_deaths = np.random.binomial(1, pd, num_time_steps) *
25    rand_life_attacks
26
27
28    # ELS Random Scenario:
29    rand_els_attacks = np.zeros(num_time_steps)
30    half_num_time_steps = int((num_time_steps)/2)
31    els_indices = np.random.choice(half_num_time_steps, num_attacks,
32    replace=False)
33    #els_indices = np.random.randint(0, half_num_time_steps, size=
34    num_attacks)
35
36    # Input attacks (1s) in random chosen timesteps
37    rand_els_attacks[els_indices] = 1
38    rand_els_deaths = np.random.binomial(1, pd, num_time_steps) *
39    rand_els_attacks
40
41
42    return rand_life_attacks, rand_life_deaths, rand_els_attacks,
43    rand_els_deaths
44
45 # Now put this into fit_bayesian_model(attacks, deaths, num_walkers,
46    num_steps, burn_in, step_size)

```

```

35 scenarios = sim_scenarios(pd = 0.2, num_time_steps = 700,
    num_attacks = 5)
36
37 life_attack_seq, life_death_seq = scenarios[0], scenarios[1]
38 els_attack_seq, els_death_seq = scenarios[2], scenarios[3]
39
40 num_walkers = 30
41 num_steps = 60
42 burn_in = 0.3
43 step_size = 2
44
45 life_bayesian_fit = fit_bayesian_model(life_attack_seq,
    life_death_seq, num_walkers, num_steps, burn_in, step_size)
46 els_bayesian_fit = fit_bayesian_model(els_attack_seq, els_death_seq,
    num_walkers, num_steps, burn_in, step_size)
47
48 # Plug these into variance_over_time(samples) to get FIG 3B
49 life_variance = variance_over_time(life_bayesian_fit, colour = '
    Blues', x_offset = 0.015)
50 els_variance = variance_over_time(els_bayesian_fit, colour = 'Reds',
    x_offset = 0)

```

B.8 Bayesian Model - Clustered ELS Attacks

```

1 def els_clustered(pd, num_time_steps, num_attacks, sims):
2     '''
3     Simulates attack and death sequences for two scenarios, attacks
4     are clustered randomly or uniformly.
5     These are then put into variance over time plot to show how
6     threat prediction varies. Results are averaged across
7     "sims" number of different runs (different lifetime and els
8     attack sequences).
9
10    :param prob_attack: Probability of attack
11    :param prob_death: Probability of death given attack happens
12    :param num_time_steps: Total number of trials or timesteps where
13    attack can occur
14
15    :return: attacks, deaths which are arrays corresponding to
16    sequences of attacks and deaths (0 is no attack/no death,
17    1 is attack/death)

```

```

13     '''
14     life_sims = np.zeros((sims, 1260, 2)) # burnin: 0.1->1620,
0.2->1440, 0.3->1260
15     els_sims = np.zeros((sims, 1260, 2))
16
17     for s in range(sims):
18         # Lifetime Random Scenario:
19         rand_life_attacks = np.zeros(num_time_steps)
20         indices = np.random.choice(num_time_steps, num_attacks,
replace=False)
21         print(indices)
22         # Input attacks (1s) in random chosen timesteps
23         rand_life_attacks[indices] = 1
24         rand_life_deaths = np.random.binomial(1, pd, num_time_steps)
* rand_life_attacks
25
26
27         # ELS clustered Scenario: random clusters of 3 attacks
simultaneously (over 3 timesteps)
28         clus_els_attacks = np.zeros(num_time_steps)
29         half_num_time_steps = int(num_time_steps/2)
30         #indices = np.random.choice(half_num_time_steps, num_attacks
, replace=False) # RANDOM SPACING
31
32         indices = np.linspace(0, half_num_time_steps, num_attacks,
dtype=int) # EQUAL SPACING
33         print(indices)
34
35         # Input attacks (1s) in random chosen timesteps
36         clus_els_attacks[indices] = 1
37         next_ind = [x+1 for x in indices]
38         print(next_ind)
39         clus_els_attacks[next_ind] = 1
40         next_ind = [x+1 for x in next_ind]
41         print(next_ind)
42         clus_els_attacks[next_ind] = 1
43
44         #print(next_ind)
45
46         rand_els_deaths = np.random.binomial(1, pd, num_time_steps)
* clus_els_attacks
47         #print(np.sum(clus_els_attacks))

```



```
48     #print(np.sum(rand_life_attacks))
49
50     num_walkers = 30
51     num_steps = 60
52     burn_in = 0.3
53     step_size = 2
54
55     life_bayesian_fit = fit_bayesian_model(rand_life_attacks,
56     rand_life_deaths, num_walkers, num_steps, burn_in, step_size)
57     els_bayesian_fit = fit_bayesian_model(clus_els_attacks,
58     rand_els_deaths, num_walkers, num_steps, burn_in, step_size)
59
60     life_sims[s,:,:] = life_bayesian_fit
61     els_sims[s,:,:] = els_bayesian_fit
62
63     # Average over sims for each step
64     life_means = np.mean(life_sims, axis=0)
65     els_means = np.mean(els_sims, axis=0)
66     #print(np.shape(life_sims))
67     #print(np.shape(els_sims))
68     # Generating plots for avg variance over time
69     life_variance = variance_over_time(life_means, colour='Blues',
70     x_offset=0.015)
71     els_variance = variance_over_time(els_means, colour='Reds',
72     x_offset=0)
73
74     return life_sims, els_sims
75
76 life_sims, els_sims = els_clustered(pd = 0.01, num_time_steps = 700,
77     num_attacks = 5, sims=1)
```

Appendix C

TD Model Python Code

C.1 Preprocessing SEFL Data

```
1 import numpy as np
2 from sklearn.preprocessing import MinMaxScaler
3 from scipy.optimize import minimize
4 from scipy.optimize import differential_evolution
5 import pandas as pd
6 import scipy.stats
7
8
9 import scipy.io as sio
10 mat_contents = sio.loadmat('Supp_Mat')
11
12 ### EXTRACTING ALL DETAILS ###
13
14 animals = mat_contents["sefl_behavior_day1"]["animal"].reshape(-1)
15 stress_type_index = mat_contents["sefl_behavior_day1"]["stress"].
    reshape(-1)
16 day1_freezing_ts = mat_contents["sefl_behavior_day1"]["
    freezing_time_series"].reshape(-1)
17 day1_smoothed_freezing = mat_contents["sefl_behavior_day1"]["
    smoothed_freezing"].reshape(-1)
18 day6_animal_index = mat_contents["sefl_behavior_day1"]["day6_index"
    ].reshape(-1)
19 day7_animal_index = mat_contents["sefl_behavior_day1"]["day7_index"
    ].reshape(-1)
20
21
```

```

22 # Storing all details:
23 details_df = pd.DataFrame({
24     "animal": animals,
25     "stress": stress_type_index,
26     "day1_freezing_time_series": day1_freezing_ts,
27     "day1_smoothed_freezing": day1_smoothed_freezing,
28     "day6_index": day6_animal_index,
29     "day7_index": day7_animal_index
30 })
31
32 # Remove invalid rows:
33 # Mouse 3 removed - not in day7 file - G11
34 # Mouse 9 removed - no data - G17
35 # Mouse 25 removed - no data in day6 file - G33
36 # Mouse 28 removed - too many shocks - G36
37 # Mouse 35 removed - too many shocks - G9
38
39 indexes_to_remove = [2, 8, 24, 27, 34]
40 details_df = details_df.drop(indexes_to_remove)
41 # Rest row indexes are removing
42 details_df = details_df.reset_index(drop=True)
43
44 # Adjusting df to be nicer to work with:
45 for i in range(len(details_df["animal"])):
46     details_df["animal"][i] = details_df["animal"][i][0]
47     details_df["stress"][i] = details_df["stress"][i][0][0]
48     details_df["day6_index"][i] = details_df["day6_index"][i][0][0]
49     details_df["day7_index"][i] = details_df["day7_index"][i][0][0]
50     details_df["day1_freezing_time_series"][i] = details_df["
51     day1_freezing_time_series"][i].tolist()
52     details_df["day1_smoothed_freezing"][i] = details_df["
53     day1_smoothed_freezing"][i].tolist()
54
55 # Changing lower case g to upper case to match day6 and day7 animal
56     names:
57 details_df["animal"] = details_df["animal"].str.replace(r'^g', 'G')
58
59 # Day 1 Shocks:
60 day1_shock_times = mat_contents["sefl_behavior_day1"]["shock_times"
61     ].reshape(-1)
62
63 # Removing rows of relevant indexes:
64 day1_shock_times = np.delete(day1_shock_times, indexes_to_remove)

```

```

60
61
62 # Now creating the actual attack sequences (0s and 1s every timestep
    ) - 33 mice
63 day1_attack_sequences = np.zeros((33, 162000))
64
65 for m in range(33):
66     times = day1_shock_times[m]
67     # Adjust from 1-based indexing
68     indexes = times-1
69
70     # Add attack (1) at each index accordingly
71     for i in indexes:
72         day1_attack_sequences[m, int(i)] = 1
73
74 # Convert each row of day1_attack_sequences into a list
75 attack_sequences_list = [row.tolist() for row in
    day1_attack_sequences]
76
77 # Create the DataFrame with each list as a row
78 day1_attack_sequences = pd.DataFrame({"attack_sequence":
    attack_sequences_list})
79 # Sanity check: check for several mice that attacks are in right
    place:
80 #np.where(np.array(day1_attack_sequences["attack_sequence"][0]) ==
    1) # Mouse G1
81
82
83 # Now attach this to details_df
84 details_df.insert(2, "day1_attack_sequences", day1_attack_sequences)
85 #details_df["day1_attack_sequences"] = day1_attack_sequences
86
87
88
89 # Extracting day6 and day7 freezing time series and smoothed
    freezing for these mice:
90 day6_animal = mat_contents["sefl_behavior_day6"]["animal"].reshape
    (-1)
91 day6_freezing_ts = mat_contents["sefl_behavior_day6"]["
    freezing_time_series"].reshape(-1)
92 day6_smoothed_freezing = mat_contents["sefl_behavior_day6"]["
    smoothed_freezing"].reshape(-1)

```

```

93
94 day6_details_df = pd.DataFrame({
95     "animal": day6_animal,
96     "day6_freezing_time_series": day6_freezing_ts,
97     "day6_smoothed_freezing": day6_smoothed_freezing
98 })
99
100 # Adjusting df to be nicer to work with:
101 for i in range(len(day6_details_df["animal"])):
102     day6_details_df["animal"][i] = day6_details_df["animal"][i
103 ] [0] [0] [0]
104     day6_details_df["day6_freezing_time_series"][i] =
105     day6_details_df["day6_freezing_time_series"][i].tolist()
106     day6_details_df["day6_smoothed_freezing"][i] = day6_details_df["
107     day6_smoothed_freezing"][i].tolist()
108
109 # Day7
110 day7_animal = mat_contents["sefl_behavior_day7"]["animal"].reshape
111 (-1)
112 day7_freezing_ts = mat_contents["sefl_behavior_day7"]["
113     freezing_time_series"].reshape(-1)
114 day7_smoothed_freezing = mat_contents["sefl_behavior_day7"]["
115     smoothed_freezing"].reshape(-1)
116
117 day7_details_df = pd.DataFrame({
118     "animal": day7_animal,
119     "day7_freezing_time_series": day7_freezing_ts,
120     "day7_smoothed_freezing": day7_smoothed_freezing
121 })
122
123 # Adjusting df to be nicer to work with:
124 for i in range(len(day7_details_df["animal"])):
125     day7_details_df["animal"][i] = day7_details_df["animal"][i
126 ] [0] [0] [0]
127     day7_details_df["day7_freezing_time_series"][i] =
128     day7_details_df["day7_freezing_time_series"][i].tolist()
129     day7_details_df["day7_smoothed_freezing"][i] = day7_details_df["
130     day7_smoothed_freezing"][i].tolist()
131
132 # Now joining day6 and day7 details into main details_df:
133 # Merging day6_details_df onto details_df based on "animal" column

```

```

126 details_df = pd.merge(details_df, day6_details_df[["animal", "
    day6_freezing_time_series"]], on="animal", how="left")
127 details_df = pd.merge(details_df, day6_details_df[["animal", "
    day6_smoothed_freezing"]], on="animal", how="left")
128
129
130 # Merging day7_details_df onto details_df based on "animal" column
131 details_df = pd.merge(details_df, day7_details_df[["animal", "
    day7_freezing_time_series"]], on="animal", how="left")
132 details_df = pd.merge(details_df, day7_details_df[["animal", "
    day7_smoothed_freezing"]], on="animal", how="left")
133
134 # Now details_df has all details I require, I can fit data for all
    days for the correct mice:
135 stressed_mice_details = details_df[details_df["stress"] == 1].copy()
136 stressed_mice_details = stressed_mice_details.reset_index(drop=True)
137
138 unstressed_mice_details = details_df[details_df["stress"] == 0].copy
    ()
139 unstressed_mice_details = unstressed_mice_details.reset_index(drop=
    True)

```

C.2 TD Model - Smoothed Freezing & Fitting SEFL Data

```

1 def TD_model_init(init, u, alpha, gammal):
2     '''
3     Evaluates TD model threat at each timestep
4
5     :param u: Input (sequence of unconditioned stimuli, i.e.
    sequences of footshocks, 0s and 1s), for different
6     contexts, i.e A and B (day 1 is context A, day 6 and 7 is
    context B).
7     :param alpha: Learning rate
8     :param gammal: Decay rate for threat
9
10    :return: TD model threat estimation levels over all timesteps
11    '''
12    T = np.zeros_like(u)
13    T[0] = init
14    for t in range(1, len(T)):
15        T[t] = T[t-1] + alpha * (u[t] - gammal * T[t-1])

```

```

16
17     # Scaling between 0.1 and 0.9
18     scaler = MinMaxScaler(feature_range=(0.1, 0.9))
19     T = scaler.fit_transform(T.reshape(-1, 1)).flatten()
20
21     return T
22
23 def NLL_td(params, init, stimuli, threats):
24     '''
25     Calculates NLL for one set of parameters given a sequence of
26     attacks and threat predictions for the TD model.
27     The likelihood calculation uses the PMF of a Bernoulli
28     distribution to calculate the log probability of observing
29     the given threat probability "threat" based on the current value
30     of T.
31
32     :param params: Array of parameters alpha and gammal
33     :param init: initialisation for threat sequence (eg end of day 1
34     context A for day 6 context A
35     :param stimuli: Sequence of shocks (attacks)
36     :param threats: Sequence of threat predictions
37
38     :return: Negative Log-Likelihood of one set of parameters.
39     '''
40     alpha, gammal = params
41     log_likelihood = 0.0
42
43     T = TD_model_init(init = init, u = stimuli, alpha = alpha,
44                       gammal = gammal)
45
46     eps = 1e-10
47     T = np.clip(T, 0+eps, 1-eps)
48
49     for t in range(len(stimuli)):
50         # Assuming threat distribution is Bernoulli - this appears
51         # to work as needed
52         # Uses probability mass function (PMF) of a Bernoulli
53         # distribution: considers threats as probabilities, i.e.
54         # each 'threat' represents the prob of observing a shock.
55         log_likelihood += threats[t] * np.log(T[t]) + (1 - threats[t
56 ]) * np.log(1 - T[t])

```

```
50     nll = -log_likelihood
51
52     return nll
53
54 def smoothing(raw_time_series, window_size):
55     smoothed_time_series = []
56     half_window = window_size // 2
57
58     # Pad the time series
59     padded_time_series = np.pad(raw_time_series, (half_window,
60     half_window), mode='edge')
61
62     for i in range(len(raw_time_series)):
63         window_values = padded_time_series[i : i + window_size]
64         smoothed_value = np.mean(window_values) # Average of window
65         smoothed_time_series.append(smoothed_value)
66
67     return np.array(smoothed_time_series)
68
69
70
71
72 ### FITTING CONTEXT A DAY 1 ### STRESSED
73
74 # Storing MANUAL smoothed freezing data for each mouse
75 day1_stressed_smoothed_freezing = np.zeros((len(
76     stressed_mice_details), len(stressed_mice_details["
77     day1_freezing_time_series"][0])))
78
79 # Applying smoothing function to each mouse freezing time series
80 for i in range(len(stressed_mice_details)):
81     raw_time_series = stressed_mice_details["
82     day1_freezing_time_series"][i]
83     smoothed_time_series = smoothing(raw_time_series, window_size =
84     15)
85     day1_stressed_smoothed_freezing[i, :] = smoothed_time_series
86
87 day1_stressed_freezing = day1_stressed_smoothed_freezing
88     # (15, 162000) use as prediction input
89 day1_stressed_attacks = stressed_mice_details["day1_attack_sequences
90     "]
```



```

85 day1_stressed_attacks = np.array(day1_stressed_attacks.to_list())
      # (15, 162000) use as stimuli input
86 initial = 0
87
88 initial_vals = np.array((0.5, 0.95))
89 bounds = [(0.05,0.9), (0.9,1)]
90
91 # Minimization for MICE
92 for m in range(0, len(stressed_mice_details)):
93     # Compute minimization for each participant
94     result = differential_evolution(NLL_td, x0=initial_vals, bounds=
bounds,
95                                     args=(initial, day1_stressed_attacks[m:m
+ 1].reshape(-1), day1_stressed_freezing[m:m + 1].reshape(-1)),
96                                     polish=True)
97
98     print(f'n_iter: {result.nit} - success: {result.success} - nll {
result.fun}')
99
100    # Store in results dataframe:
101    TD_stressed_results["Day 1 Context A NLL"][m] = result.fun
102    alpha, gamma1 = result.x
103    TD_stressed_results["Day 1 Context A Params"][m] = [alpha,
gamma1]
104
105
106
107 ### FITTING CONTEXT B DAY 1 ### STRESSED
108 day1_stressed_attacks_B = np.zeros_like(day1_stressed_attacks)
      # (15, 162000) use as stimuli input
109 initial = 0
110
111 # Minimization for MICE
112 for m in range(0, len(stressed_mice_details)):
113     # Compute TD threat fit from parameters for this day, compute
NLL
114     a, g = TD_stressed_results["Day 1 Context A Params"][m]
115     day1_stressed_freezing_B = TD_model_init(init=initial, u=
day1_stressed_attacks_B[m,:], alpha= a, gamma1=g) # (15, 162000)
      use as threat input
116
117     params = [a, g]

```

```

118     nll = NLL_td(params=params, init=initial, stimuli=
day1_stressed_attacks_B[m,:], threats=day1_stressed_freezing_B)
119     # Store in results dataframe:
120     TD_stressed_results["Day 1 Context B NLL"][m] = nll
121
122
123
124
125
126 ### FITTING CONTEXT B DAY 6 ### STRESSED
127 # Storing MANUAL smoothed freezing data for each mouse
128 day6_stressed_smoothed_freezing = np.zeros((len(
    stressed_mice_details), len(stressed_mice_details["
    day6_freezing_time_series"][0])))
129
130 # Applying smoothing function to each mouse freezing time series
131 for i in range(len(stressed_mice_details)):
132     raw_time_series = stressed_mice_details["
    day6_freezing_time_series"][i]
133     smoothed_time_series = smoothing(raw_time_series, window_size =
    15)
134     day6_stressed_smoothed_freezing[i, :] = smoothed_time_series
135
136 day6_stressed_freezing = day6_stressed_smoothed_freezing
    # (15, 18000) use as prediction input
137 day6_stressed_attacks_B = np.zeros_like(day6_stressed_freezing)
    # (15, 18000) use as stimuli input
138
139 for m in range(len(day6_stressed_attacks_B)):
140     day6_stressed_attacks_B[m,8999] = 1
141     initial = 0
142
143     initial_vals = np.array((0.5, 0.95))
144     bounds = [(0.05,0.9), (0.9,1)]
145     # Minimization for MICE
146     for m in range(0, len(stressed_mice_details)):
147         # Compute minimization for each participant
148         result = differential_evolution(NLL_td, x0=initial_vals, bounds=
    bounds,
149
    args=(initial, day6_stressed_attacks_B[m:
    m + 1].reshape(-1), day6_stressed_freezing[m:m + 1].reshape(-1)),
150
    polish=True)

```

```

151
152     print(f'n_iter: {result.nit} - success: {result.success} - nll {
result.fun}')
153
154     # Store in results dataframe:
155     TD_stressed_results["Day 6 Context B NLL"][m] = result.fun
156     alpha, gammal = result.x
157     TD_stressed_results["Day 6 Context B Params"][m] = [alpha,
gamma1]
158
159
160
161
162     ### FITTING CONTEXT A DAY 6 ### STRESSED
163     day6_stressed_attacks_A = np.zeros_like(day6_stressed_attacks_B)
        # (15, 162000) use as stimuli input
164
165     # Minimization for MICE
166     for m in range(0, len(stressed_mice_details)):
167         # Finding initial for each mouse
168         a_prev, g_prev = TD_stressed_results["Day 1 Context A Params"][m
]
169         initial = TD_model_init(init=0, u=day1_stressed_attacks[m,:],
alpha=a_prev, gammal=g_prev)[-1]
170
171         # Compute TD threat fit from parameters for this day, compute
NLL
172         a, g = TD_stressed_results["Day 6 Context B Params"][m]
173         day6_stressed_freezing_A = TD_model_init(init=initial, u=
day6_stressed_attacks_A[m:], alpha= a, gammal=g) # (15, 162000)
use as threat input
174
175         params = [a, g]
176         nll = NLL_td(params=params, init=initial, stimuli=
day6_stressed_attacks_A[m:], threats=day6_stressed_freezing_A)
177         # Store in results dataframe:
178         TD_stressed_results["Day 6 Context A NLL"][m] = nll
179
180
181
182
183

```



```

215     print(f'n_iter: {result.nit} - success: {result.success} - nll {
result.fun}')
216
217     # Store in results dataframe:
218     TD_stressed_results["Day 7 Context B NLL"][m] = result.fun
219     alpha, gamma1 = result.x
220     TD_stressed_results["Day 7 Context A Params"][m] = [alpha,
gamma1]
221
222
223
224
225     ### FITTING CONTEXT A DAY 7 ### STRESSED
226     day7_stressed_attacks_A = np.zeros_like(day7_stressed_attacks_B)
227         # (15, 18000) use as stimuli input
228
229     # Minimization for MICE
230     for m in range(0, len(stressed_mice_details)):
231         # Finding initial for each mouse
232         a_prev_prev, g_prev_prev = TD_stressed_results["Day 1 Context A
Params"][m]
233         initial_prev = TD_model_init(init=0, u=day1_stressed_attacks[m
:], alpha=a_prev_prev, gamma1=g_prev_prev)[-1]
234
235         a_prev, g_prev = TD_stressed_results["Day 6 Context B Params"][m
]
236         initial = TD_model_init(init=initial_prev, u=
day6_stressed_attacks_A[m:], alpha=a_prev, gamma1=g_prev)[-1]
237
238         # Compute TD threat fit from parameters for this day, compute
NLL
239         a, g = TD_stressed_results["Day 7 Context A Params"][m]
240         day7_stressed_freezing_A = TD_model_init(init=initial, u=
day7_stressed_attacks_A[m:], alpha= a, gamma1=g) # (15, 162000)
241         use as threat input
242
243         params = [a, g]
244         nll = NLL_td(params=params, init=initial, stimuli=
day7_stressed_attacks_A[m:], threats=day7_stressed_freezing_A)
245         # Store in results dataframe:
246         TD_stressed_results["Day 7 Context A NLL"][m] = nll

```

```
246
247
248
249
250
251 ### NOW FOR UNSTRESSED ###
252
253
254
255
256
257
258 ### FITTING CONTEXT A DAY 1 ### UNSTRESSED
259
260 # Storing MANUAL smoothed freezing data for each mouse
261 day1_unstressed_smoothed_freezing = np.zeros((len(
    unstressed_mice_details), len(unstressed_mice_details["
    day1_freezing_time_series"][0])))
262
263 # Applying smoothing function to each mouse freezing time series
264 for i in range(len(unstressed_mice_details)):
265     raw_time_series = unstressed_mice_details["
    day1_freezing_time_series"][i]
266     smoothed_time_series = smoothing(raw_time_series, window_size =
    15)
267     day1_unstressed_smoothed_freezing[i, :] = smoothed_time_series
268
269 day1_unstressed_freezing = day1_unstressed_smoothed_freezing
    # (15, 162000) use as prediction input
270 day1_unstressed_attacks = np.zeros_like(day1_unstressed_freezing)
271 initial = 0
272
273 initial_vals = np.array((0.5, 0.95))
274 bounds = [(0.05,0.9), (0.9,1)]
275
276 # Minimization for MICE
277 for m in range(0, len(unstressed_mice_details)):
278     # Compute minimization for each participant
279     result = differential_evolution(NLL_td, x0=initial_vals, bounds=
    bounds,
280                                     args=(initial, day1_unstressed_attacks[m:
    m + 1].reshape(-1), day1_unstressed_freezing[m:m + 1].reshape(-1))
```

```

),
281         polish=True)
282
283     print(f'n_iter: {result.nit} - success: {result.success} - nll {
result.fun}')
284
285     # Store in results dataframe:
286     TD_unstressed_results["Day 1 Context A NLL"][m] = result.fun
287     alpha, gamma1 = result.x
288     TD_unstressed_results["Day 1 Context A Params"][m] = [alpha,
gamma1]
289
290
291
292 ### FITTING CONTEXT B DAY 1 ### UNSTRESSED
293 day1_unstressed_attacks_B = np.zeros_like(day1_unstressed_attacks)
294     # (15, 162000) use as stimuli input
295 initial = 0
296
297 # Minimization for MICE
298 for m in range(0, len(unstressed_mice_details)):
299     # Compute TD threat fit from parameters for this day, compute
NLL
300     a, g = TD_unstressed_results["Day 1 Context A Params"][m]
301     day1_unstressed_freezing_B = TD_model_init(init=initial, u=
day1_unstressed_attacks_B[m,:], alpha= a, gamma1=g) # (15,
162000) use as threat input
302
303     params = [a, g]
304     nll = NLL_td(params=params, init=initial, stimuli=
day1_unstressed_attacks_B[m,:], threats=
day1_unstressed_freezing_B)
305     # Store in results dataframe:
306     TD_unstressed_results["Day 1 Context B NLL"][m] = nll
307
308
309
310
311 ### FITTING CONTEXT B DAY 6 ### UNSTRESSED
312 # Storing MANUAL smoothed freezing data for each mouse
313 day6_unstressed_smoothed_freezing = np.zeros((len(

```

```

    unstressed_mice_details), len(unstressed_mice_details["
    day6_freezing_time_series"][0]))
314
315 # Applying smoothing function to each mouse freezing time series
316 for i in range(len(unstressed_mice_details)):
317     raw_time_series = unstressed_mice_details["
    day6_freezing_time_series"][i]
318     smoothed_time_series = smoothing(raw_time_series, window_size =
    15)
319     day6_unstressed_smoothed_freezing[i, :] = smoothed_time_series
320
321 day6_unstressed_freezing = day6_unstressed_smoothed_freezing
    # (15, 18000) use as prediction input
322 day6_unstressed_attacks_B = np.zeros_like(day6_unstressed_freezing)
    # (15, 18000) use as stimuli input
323
324 for m in range(len(day6_unstressed_attacks_B)):
325     day6_unstressed_attacks_B[m, 8999] = 1
326 initial = 0
327
328 initial_vals = np.array((0.5, 0.95))
329 bounds = [(0.05, 0.9), (0.9, 1)]
330 # Minimization for MICE
331 for m in range(0, len(unstressed_mice_details)):
332     # Compute minimization for each participant
333     result = differential_evolution(NLL_td, x0=initial_vals, bounds=
    bounds,
334                                     args=(initial, day6_unstressed_attacks_B[
    m:m + 1].reshape(-1), day6_unstressed_freezing[m:m + 1].reshape
    (-1)),
335                                     polish=True)
336
337     print(f'n_iter: {result.nit} - success: {result.success} - nll {
    result.fun}')
338
339     # Store in results dataframe:
340     TD_unstressed_results["Day 6 Context B NLL"][m] = result.fun
341     alpha, gamma1 = result.x
342     TD_unstressed_results["Day 6 Context B Params"][m] = [alpha,
    gamma1]
343
344

```



```

345
346
347 ### FITTING CONTEXT A DAY 6 ### UNSTRESSED
348 day6_unstressed_attacks_A = np.zeros_like(day6_unstressed_attacks_B)
      # (15, 162000) use as stimuli input
349
350 # Minimization for MICE
351 for m in range(0, len(unstressed_mice_details)):
352     # Finding initial for each mouse
353     a_prev, g_prev = TD_unstressed_results["Day 1 Context A Params"
      ][m]
354     initial = TD_model_init(init=0, u=day1_unstressed_attacks[m,:],
      alpha=a_prev, gamma=g_prev)[-1]
355
356     # Compute TD threat fit from parameters for this day, compute
      NLL
357     a, g = TD_unstressed_results["Day 6 Context B Params"][m]
358     day6_unstressed_freezing_A = TD_model_init(init=initial, u=
      day6_unstressed_attacks_A[m:], alpha= a, gamma=g) # (15,
      162000) use as threat input
359
360     params = [a, g]
361     nll = NLL_td(params=params, init=initial, stimuli=
      day6_unstressed_attacks_A[m:], threats=
      day6_unstressed_freezing_A)
362     # Store in results dataframe:
363     TD_unstressed_results["Day 6 Context A NLL"][m] = nll
364
365
366
367
368
369
370 ### FITTING CONTEXT B DAY 7 ### UNSTRESSED
371 # Storing MANUAL smoothed freezing data for each mouse
372 day7_unstressed_smoothed_freezing = np.zeros((len(
      unstressed_mice_details), len(unstressed_mice_details["
      day7_freezing_time_series"][0])))
373
374 # Applying smoothing function to each mouse freezing time series
375 for i in range(len(unstressed_mice_details)):
376     raw_time_series = unstressed_mice_details["

```

```

    day7_freezing_time_series"][i]
377     smoothed_time_series = smoothing(raw_time_series, window_size =
    15)
378     day7_unstressed_smoothed_freezing[i, :] = smoothed_time_series
379
380
381 day7_unstressed_freezing = day7_unstressed_smoothed_freezing
    # (15, 18000) use as prediction input
382 day7_unstressed_attacks_B = np.zeros_like(day7_unstressed_freezing)
    # (15, 18000) use as stimuli input
383
384 initial_vals = np.array((0.5, 0.95))
385 bounds = [(0.05,0.9), (0.9,1)]
386 # Minimization for MICE
387 for m in range(0, len(unstressed_mice_details)):
388     # Getting initial for each mouse:
389     a_prev_prev, g_prev_prev = TD_unstressed_results["Day 1 Context
    A Params"][m]
390     initial_prev = TD_model_init(init=0, u=day1_unstressed_attacks_B
    [m,:], alpha=a_prev_prev, gamma1=g_prev_prev)[-1]
391
392     a, g = TD_unstressed_results["Day 6 Context B Params"][m]
393     initial = TD_model_init(init=initial_prev, u=
    day6_unstressed_attacks_B[m,:], alpha=a, gamma1=g)[-1]
394
395     # Compute minimization for each participant
396     result = differential_evolution(NLL_td, x0=initial_vals, bounds=
    bounds,
397                                     args=(initial, day7_unstressed_attacks_B[
    m:m + 1].reshape(-1), day7_unstressed_freezing[m:m + 1].reshape
    (-1)),
398                                     polish=True)
399
400     print(f'n_iter: {result.nit} - success: {result.success} - nll {
    result.fun}')
401
402     # Store in results dataframe:
403     TD_unstressed_results["Day 7 Context B NLL"][m] = result.fun
404     alpha, gamma1 = result.x
405     TD_unstressed_results["Day 7 Context A Params"][m] = [alpha,
    gamma1]
406

```

```

407
408
409
410 ### FITTING CONTEXT A DAY 7 ### UNSTRESSED
411 day7_unstressed_attacks_A = np.zeros_like(day7_unstressed_attacks_B)
      # (15, 18000) use as stimuli input
412
413 # Minimization for MICE
414 for m in range(0, len(unstressed_mice_details)):
415     # Finding initial for each mouse
416     a_prev_prev, g_prev_prev = TD_unstressed_results["Day 1 Context
      A Params"][m]
417     initial_prev = TD_model_init(init=0, u=day1_unstressed_attacks[m
      ,:], alpha=a_prev_prev, gamma1=g_prev_prev)[-1]
418
419     a_prev, g_prev = TD_unstressed_results["Day 6 Context B Params"
      ][m]
420     initial = TD_model_init(init=initial_prev, u=
      day6_unstressed_attacks_A[m,:], alpha=a_prev, gamma1=g_prev)[-1]
421
422     # Compute TD threat fit from parameters for this day, compute
      NLL
423     a, g = TD_unstressed_results["Day 7 Context A Params"][m]
424     day7_unstressed_freezing_A = TD_model_init(init=initial, u=
      day7_unstressed_attacks_A[m,:], alpha= a, gamma1=g) # (15,
      162000) use as threat input
425
426     params = [a, g]
427     nll = NLL_td(params=params, init=initial, stimuli=
      day7_unstressed_attacks_A[m,:], threats=
      day7_unstressed_freezing_A)
428     # Store in results dataframe:
429     TD_unstressed_results["Day 7 Context A NLL"][m] = nll

```

C.3 TD Model - Computing BIC Scores

```

1 ### BIC SCORES FOR TD MODEL ###
2
3 # BIC = 2 * NLL + p*log(n)
4

```

```
5 # where p = number of params, n = number of observations (162000
   day1 and 18000 for day 6 and 7)
6 p = 2
7 # Stressed BIC
8 TD_stressed_results["BIC Score"] = None
9 for m in range(len(TD_stressed_results)):
10     Day1_NLL = (TD_stressed_results["Day 1 Context A NLL"][m] +
11                TD_stressed_results["Day 1 Context B NLL"][m])
12
13     Day6_NLL = (TD_stressed_results["Day 6 Context A NLL"][m] +
14                TD_stressed_results["Day 6 Context B NLL"][m])
15
16     Day7_NLL = (TD_stressed_results["Day 7 Context A NLL"][m] +
17                TD_stressed_results["Day 7 Context B NLL"][m])
18
19     Day1_BIC = 2 * Day1_NLL + p * np.log(162000)
20
21     Day6_BIC = 2 * Day6_NLL + p * np.log(18000)
22
23     Day7_BIC = 2 * Day7_NLL + p * np.log(18000)
24
25     TD_stressed_results["BIC Score"][m] = Day1_BIC + Day6_BIC +
   Day7_BIC
26
27 # Unstressed BIC
28 TD_unstressed_results["BIC Score"] = None
29 for m in range(len(TD_unstressed_results)):
30     Day1_NLL = (TD_unstressed_results["Day 1 Context A NLL"][m] +
31                TD_unstressed_results["Day 1 Context B NLL"][m])
32
33     Day6_NLL = (TD_unstressed_results["Day 6 Context A NLL"][m] +
34                TD_unstressed_results["Day 6 Context B NLL"][m])
35
36     Day7_NLL = (TD_unstressed_results["Day 7 Context A NLL"][m] +
37                TD_unstressed_results["Day 7 Context B NLL"][m])
38
39     Day1_BIC = 2 * Day1_NLL + p * np.log(162000)
40
41     Day6_BIC = 2 * Day6_NLL + p * np.log(18000)
42
43     Day7_BIC = 2 * Day7_NLL + p * np.log(18000)
44
```

```

45 TD_unstressed_results["BIC Score"][m] = Day1_BIC + Day6_BIC +
    Day7_BIC

```

C.4 TD Model - Multivariate Sampled Parameter Recovery

```

1 import matplotlib.pyplot as plt
2 import numpy as np
3 from scipy.optimize import minimize
4 import scipy
5 import pandas as pd
6 from scipy.optimize import differential_evolution
7
8
9 # Model Param lists (taken from paper supplementary material)
10 # Loading matlab file from paper - creates dict "mat_contents" which
    stores all variables
11 import scipy.io as sio
12 mat_contents = sio.loadmat('Supp_Mat')
13
14
15 def day_1(shocks):
16     '''
17     Evaluates day1 of SEFL experiment.
18
19     :return: Outputs sequence of 15 shocks randomly over 90 mins
    (5400 secs)
20     '''
21     min_shock_int = 4 * 1800
22     max_shock_int = 6 * 1800
23     # 90 mins = 5400 secs
24     day_1 = np.zeros(162000)
25
26     all_shock_intervals = []
27     while np.sum(day_1) < shocks:
28         # Generate random interval between 4 and 8 mins
29         shock_interval = np.random.randint(min_shock_int,
    max_shock_int)
30         # Add this to the list of all interval times
31         all_shock_intervals.append(shock_interval)

```

```

32     # Calculate the shock index by summing cumulative intervals
33     shock_index = np.sum(all_shock_intervals)
34     # Add shock to day_1 array
35     day_1[shock_index] = 1
36     return day_1
37
38 def NLL_td(params, stimuli, threats):
39     '''
40     Calculates NLL for one set of parameters given a sequence of
41     attacks and threat predictions for the TD model.
42     The likelihood calculation uses the PMF of a Bernoulli
43     distribution to calculate the log probability of observing
44     the given threat probability "threat" based on the current value
45     of T.
46
47     :param params: Array of parameters alpha and gamma1
48     :param stimuli: Sequence of shocks (attacks)
49     :param threats: Sequence of threat predictions
50
51     :return: Negative Log-Likelihood of one set of parameters.
52     '''
53     alpha, gamma1 = params
54     T_prev = 0.0
55     log_likelihood = 0.0
56
57     for u, threat in zip(stimuli, threats):
58
59         T = T_prev + alpha * (u - gamma1 * T_prev)
60
61         # Getting some invalid log values, divisions by 0 so adding
62         # small value to stop this
63         eps = 1e-10
64         T = np.clip(T, 0+eps, 1-eps)
65
66         # Assuming threat distribution is Bernoulli - this appears
67         # to work as needed
68         # Uses probability mass function (PMF) of a Bernoulli
69         # distribution: considers threats as probabilities, i.e.
70         # each 'threat' represents the prob of observing a shock.
71         log_likelihood += np.log((T ** threat) * ((1 - T) ** (1 -
72         threat)))
73

```

```

67     T_prev = T
68
69     nll = -log_likelihood
70
71     return nll
72
73
74 def TD_model(u, alpha, gamma1):
75     '''
76     Evaluates TD model threat at each timestep
77
78     :param u: Input (sequence of unconditioned stimuli, i.e.
79     sequences of footshocks, 0s and 1s), for different
80     contexts, i.e A and B (day 1 is context A, day 6 and 7 is
81     context B).
82
83     :param alpha: Learning rate
84     :param gamma1: Decay rate for threat
85
86     :return: TD model threat estimation levels over all timesteps
87     '''
88     T = np.zeros_like(u)
89     for t in range(1, len(T)):
90         T[t] = T[t-1] + alpha * (u[t] - gamma1 * T[t-1])
91
92     return T
93
94 # Use this now to generate 100 simulations and fit params
95 # Draw 100 sets of param values from suitable dist, e.g.
96     multivariate normal distribution
97 # Set covariance to 0, choose small numbers for variance.
98 # Check values before using. Clip to 0 and 1.
99
100 alpha_mean = 0.5
101 gamma1_mean = 0.95
102 means = [alpha_mean, gamma1_mean]
103
104 alpha_var = 0.03
105 gamma1_var = 0.001
106 covs = np.zeros((2,2))
107 covs[0,0] = alpha_var
108 covs[1,1] = gamma1_var
109

```

```

106 sample_params = np.random.multivariate_normal(mean=means, cov=covs,
        size=100)
107 sample_params[:,0] = np.clip(sample_params[:,0], 0, 1)
108 sample_params[:,1] = np.clip(sample_params[:,1], 0.9, 1)
109
110
111 '''
112 # Plot showing parameter values for reference
113 fig, ax = plt.subplots(2)
114 fig.suptitle("Sampled Parameter Values")
115 # add a big axes, hide frame
116 fig.add_subplot(111, frameon=False)
117 # hide tick and tick label of the big axes
118 plt.tick_params(labelcolor='none', top=False, bottom=False, left=
        False, right=False)
119 plt.grid(False)
120 x = np.linspace(0, 100, 100)
121 ax[0].plot(x, sample_params[:,0], "r+", label="Sampled $alpha$")
122 ax[1].plot(x, sample_params[:,1], "bo", label="Sampled $gamma_1$")
123 plt.xlabel("Simulation (Individual Mouse)")
124 plt.ylabel("Sampled Parameter Value")
125 ax[1].legend(fontsize="8")
126 ax[0].legend(fontsize="8")
127 plt.show()
128 '''
129
130 # Now create simulations of attacks and corresponding threat
        predictions for each pair of params:
131 sampled_attacks = np.zeros((100, 162000))
132 sampled_predictions = np.zeros((100, 162000))
133
134 for index, (a, g1) in enumerate(zip(sample_params[:,0],
        sample_params[:,1])):
135     # Generate day 1 sequence of attacks for each
136     attacks = day_1(15)
137     # Generate threat predictions from this simulated sequence of
        attacks
138     predictions = TD_model(u = attacks, alpha = a, gammal = g1)
139
140     # Add to sampled attacks/preds arrays
141     sampled_attacks[index, :] = attacks
142     sampled_predictions[index, :] = predictions

```


C.5 TD Model - List Sampled Parameter Recovery

```
1 # Extracting lists form supplementary information
2 alpha_list = mat_contents["learning_rate_list"][0].tolist()
3 gamma1_list = mat_contents["decay_list"][0].tolist()
4
5 sample_params = np.zeros((100, 2))
6
7 for s in range(100):
8     sample_params[s, 0] = np.random.choice(alpha_list)
9     sample_params[s, 1] = np.random.choice(gamma1_list)
10
11
12 # Now create simulations of attacks and corresponding threat
13     predictions for each set of params
14
15 sampled_attacks = np.zeros((100, 162000))
16 sampled_predictions = np.zeros((100, 162000))
17
18 for index, (a, g1) in enumerate(zip(sample_params[:,0],
19     sample_params[:,1])):
20     # Generate day 1 sequence of attacks for each
21     attacks = day_1(15)
22     # Generate threat predictions from this simulated sequence of
23     attacks
24     predictions = TD_model(u = attacks, alpha = a, gamma1 = g1)
25
26     # Add to sampled attacks/preds arrays
27     sampled_attacks[index, :] = attacks
28     sampled_predictions[index, :] = predictions
29
30 # Sanity check
31 np.sum(sampled_predictions, axis=1)
32
33 # Fitting parameter values to these simulations (see if we can
34     recover params)
35
36 initial_vals = np.array((0.5, 0.95))
37
38 bounds = [(0.05, 0.9), (0.9, 1)]
39
40 opt_sampled_data = np.zeros((len(sampled_predictions[:,0]), 2))
```


Appendix D

TD-Momentum Model Code

D.1 TD-Momentum Model - Example of SEFL Model Fit & Related Freezing (Fig 3.7)

```
1 import numpy as np
2 from scipy.optimize import minimize
3 import scipy.io as sio
4 mat_contents = sio.loadmat('Supp_Mat')
5 import matplotlib.pyplot as plt
6
7
8 ### PRE-PROCESSING START ###
9
10 # DAY1: 162000 steps, 38 mice
11 # Day 1 Freezing
12 # This is the first mice freezing![:,1] for next etc, 38 total.
    Shape to be (38, 162000), 162000 timesteps
13
14 day1_freezing = np.zeros((38, 162000))
15
16 for m in range(38):
17     day1_freezing[m,:] = mat_contents["sefl_behavior_day1"]["
        smoothed_freezing"][:,m][0].reshape(-1)
18
19 # Remove rows for removed mice
20 day1_freezing = np.delete(day1_freezing, 8, axis=0)
21 day1_freezing = np.delete(day1_freezing, 33, axis=0)
22 day1_freezing = np.delete(day1_freezing, 26, axis=0)
```

```

23
24
25 # Day 1 Shocks:
26 # 38 total
27 # Mouse 9 removed - no data
28 # Mouse 28 removed - too many shocks
29 # Mouse 35 removed - too many shocks
30 day1_shock_times = mat_contents["sefl_behavior_day1"]["shock_times"]
31 # Remove indexes from day1_shock_times
32 day1_shock_times_modified = np.delete(day1_shock_times, 8, axis=1)
33 day1_shock_times_modified = np.delete(day1_shock_times_modified, 33,
    axis=1)
34 day1_shock_times_modified = np.delete(day1_shock_times_modified, 26,
    axis=1)
35
36 day1_shock_times = np.zeros((35, 15))
37
38 for m in range(35):
39     day1_shock_times[m,:] = day1_shock_times_modified[:,m][0].
    reshape(-1)
40
41 # Now creating the actual attack sequences (0s and 1s every timestep
    )
42 day1_attack_sequences = np.zeros((35, 162000))
43
44 for m in range(35):
45     times = day1_shock_times[m,:]
46     indexes = times-1
47
48     for i in indexes:
49         day1_attack_sequences[m, int(i)] = 1
50
51
52 # Array of stress type for each mouse, 0 is unstressed, 1 is
    stressed (unstressed are controls = no shocks on day1)
53 day1_stress_type = np.zeros(38)
54
55 for m in range(38):
56     day1_stress_type[m] = mat_contents["sefl_behavior_day1"]["stress
    "][0][m][0][0]
57
58 # Removing rows for removed mice

```

```

59 day1_stress_type = np.delete(day1_stress_type, 8, axis=0)
60 day1_stress_type = np.delete(day1_stress_type, 33, axis=0)
61 day1_stress_type = np.delete(day1_stress_type, 26, axis=0)
62
63 # Final arrays for stressed and unstressed mice indexes
64 stressed_indexes = np.where(day1_stress_type == 1)[0].tolist()
65 unstressed_indexes = np.where(day1_stress_type == 0)[0].tolist()
66
67
68 ### PRE-PROCESSING END ###
69
70
71
72 ### EXTRACTING EXAMPLE MOUSE FOOTSHOCKS TO FEED INTO NEW SCALED TD
    MOM MODEL TO REPRODUCE FIG5A ###
73
74
75 def TD_momentum_model_days1_6_7(init, contexts, alpha, gamma1,
    gamma2, f):
76     '''
77     Evaluates TD Momentum model threat at each timestep in each
78     context only, no changing contexts, assumes the same
79     agent in one context for the whole time, creates threat in other
80     context as momentum term only.
81
82     :param init: Provides initialisation points for threat readings
83     in both contexts (days 6 and 7)
84     :param contexts: Sequences of unconditioned stimuli across all
85     contexts. e.g. sefl(sims = 2)[1] for 2nd mouse sim
86     :param u: Key of chosen context to compute TD momentum model for
87     . (day1, day6, day7)
88     :param alpha: Learning rate
89     :param gamma1: Decay rate for threat
90     :param gamma2: Decay rate for momentum across all contexts
91     :param f: Scaling Constant for momentum term
92
93     :return: TD Momentum model threat estimation levels over all
94     timesteps
95     '''
96     # Initialise momentum array
97     m = np.zeros_like(contexts[:,0])

```

```

93     T_A = np.zeros_like(contexts[:,0])
94     T_B = np.zeros_like(contexts[:,1])
95
96     # Set initialisation
97     T_A[0] = init[0]
98     T_B[0] = init[1]
99
100    for t in range(1, len(contexts[:,0])):
101        # Set PE to 0 at every time step before computing PE for
102        # current step
103        PE = 0
104        PE += alpha * (contexts[:, 0][t] - T_A[t - 1])
105        PE += alpha * (contexts[:, 1][t] - T_B[t - 1])
106
107        m[t] = m[t-1] + gamma2 * PE
108        m[t] = np.clip(m[t], 0, 1) # NEWLY ADDED 25/07/23
109
110        T_A[t] = T_A[t - 1] + alpha * (contexts[:, 0][t] - gamma1 *
111        T_A[t - 1]) + f * m[t]
112        T_B[t] = T_B[t - 1] + alpha * (contexts[:, 1][t] - gamma1 *
113        T_B[t - 1]) + f * m[t]
114
115        T_A[t] = np.clip(T_A[t], 0, 1)
116        T_B[t] = np.clip(T_B[t], 0, 1)
117
118    return T_A, T_B
119
120    # Freezing and attacks for ALL mice (35)
121    day1_stressed_freezing = day1_freezing[stressed_indexes]
122    day1_stressed_attacks = day1_attack_sequences[stressed_indexes]
123
124    ### DAY 1 ###
125    # Taking first stressed mouse as example, extracting shock sequence
126    # for day 1: 0
127    mouse_1_day1_context_A = day1_stressed_attacks[10]
128    mouse_1_day1_context_B = np.zeros_like(mouse_1_day1_context_A)
129
130    # Day 1 threat
131    day1_contexts_stacked = np.column_stack((mouse_1_day1_context_A,
132    mouse_1_day1_context_B))
133    x = np.linspace(0, 90, len(mouse_1_day1_context_A))

```

```

130
131 init_day1 = [0, 0]
132 A_threat_day1, B_threat_day1 = TD_momentum_model_days1_6_7(init =
      init_day1, contexts = day1_contexts_stacked, alpha =0.000015,
133                                     gamma1 =
      0.9999, gamma2 = 0.01, f = 0.1)
134
135
136 ### Day 6 ###
137 # in both contexts for 10 mins, single shock in context B halfway
      through
138
139 mouse_1_day6_context_A = np.zeros(18000) #
      18000 timesteps on day 6 and day 7 (10 mins exposure)
140 mouse_1_day6_context_B = np.zeros_like(mouse_1_day6_context_A)
141 mouse_1_day6_context_B[8999] = 1 #
      Single attack at halfway point
142
143 # day 6 threat:
144 day6_contexts_stacked = np.column_stack((mouse_1_day6_context_A,
      mouse_1_day6_context_B))
145 x = np.linspace(0, 10, len(mouse_1_day6_context_A))
146
147 init_day6 = [A_threat_day1[-1], B_threat_day1[-1]]
148 A_threat_day6, B_threat_day6 = TD_momentum_model_days1_6_7(init =
      init_day6, contexts = day6_contexts_stacked, alpha =0.000015,
149                                     gamma1 =
      0.9999, gamma2 = 0.01, f = 0.1)
150
151
152 ### DAY 7 ###
153 mouse_1_day7_context_A = np.zeros(18000) #
      18000 timesteps on day 6 and day 7 (10 mins exposure)
154 mouse_1_day7_context_B = np.zeros_like(mouse_1_day7_context_A)
155
156 day7_contexts_stacked = np.column_stack((mouse_1_day7_context_A,
      mouse_1_day7_context_B))
157 x = np.linspace(0, 10, len(mouse_1_day7_context_A))
158
159 init_day7 = [A_threat_day6[-1], B_threat_day6[-1]]
160 A_threat_day7, B_threat_day7 = TD_momentum_model_days1_6_7(init =
      init_day7, contexts = day7_contexts_stacked, alpha =0.000015,

```



```

161                                     gamma1 =
    0.9999, gamma2 = 0.01, f = 0.1)
162
163
164
165 # Scale function for 6 arrays: Contexts A and B for days 1, 6 and 7
166 def combined_rescale_example(arrs, new_min, new_max):
167     combined_min = min(np.min(arr) for arr in arrs)
168     combined_max = max(np.max(arr) for arr in arrs)
169     rescaled_arrs = [(arr - combined_min) * (new_max - new_min) / (
    combined_max - combined_min) + new_min for arr in arrs]
170     return rescaled_arrs
171
172 # Scaling all days/contexts together:
173 scale_arrays = [A_threat_day1, B_threat_day1, A_threat_day6,
    B_threat_day6, A_threat_day7, B_threat_day7]
174
175 A_threat_day1, B_threat_day1, A_threat_day6, B_threat_day6,
    A_threat_day7, B_threat_day7 = combined_rescale_example(
    scale_arrays, 0.1, 0.9)
176
177
178 ### G34 FREEZING DATA ###
179 ### DAY 1 ###
180
181 mouse_1_day1_freezing = mat_contents["sefl_behavior_day1"]["
    smoothed_freezing"][:,25][0].reshape(-1)
182 mouse_1_day1_freezing = mouse_1_day1_freezing * 100
183 x = np.linspace(0, 90, len(mouse_1_day1_freezing))
184
185 ### DAY 6 ###
186 mouse_1_day6_freezing = mat_contents["sefl_behavior_day6"]["
    smoothed_freezing"][:,25][0].reshape(-1)
187 mouse_1_day6_freezing = mouse_1_day6_freezing * 100
188 x = np.linspace(0, 10, len(mouse_1_day6_freezing))
189
190 ### DAY 7 ###
191 mouse_1_day7_freezing = mat_contents["sefl_behavior_day7"]["
    smoothed_freezing"][:,24][0].reshape(-1)
192 mouse_1_day7_freezing = mouse_1_day7_freezing * 100
193 x = np.linspace(0, 10, len(mouse_1_day7_freezing))
194

```

```

195
196
197
198
199 # Plotting RL Momentum Threat example and actual smoothed freezing
200 fig, axes = plt.subplots(2, 3, figsize=(15, 10))
201
202 # Plotting TD-Momentum Threat data for each day
203 days = [1, 6, 7]
204 for i, day in enumerate(days):
205     if day == 1:
206         x = np.linspace(0, 90, len(mouse_1_day1_context_A))
207     else:
208         x = np.linspace(0, 10, len(mouse_1_day6_context_A)) if day
209         == 6 else np.linspace(0, 10, len(mouse_1_day7_context_A))
210
211     # Plot the threat data and context data on the top row axes
212     axes[1, i].plot(x, A_threat_day1 if day == 1 else A_threat_day6
213     if day == 6 else A_threat_day7, label="A")
214     axes[1, i].plot(x, B_threat_day1 if day == 1 else B_threat_day6
215     if day == 6 else B_threat_day7, label="B")
216     axes[1, i].plot(x, mouse_1_day1_context_A if day == 1 else
217     mouse_1_day6_context_A if day == 6 else mouse_1_day7_context_A,
218     label="A input", alpha=0.20)
219     axes[1, i].plot(x, mouse_1_day1_context_B if day == 1 else
220     mouse_1_day6_context_B if day == 6 else mouse_1_day7_context_B,
221     label="B input", alpha=0.20)
222     axes[1, i].set_xlabel("Time (mins)", fontsize=15)
223     axes[1, i].set_ylabel("TD-Momentum Threat", fontsize=15)
224     axes[1, i].set_ylim(0, 1)
225     axes[1, i].legend()
226
227     if i != 0:
228         axes[1, i].set_yticklabels([])
229         axes[1, i].set_ylabel("")
230
231 # Hide the x-axis labels for the top row
232 plt.setp(axes[0, :], xticks=[])
233
234 # Plotting Freezing data for each day
235 for i, day in enumerate(days):
236     if day == 1:

```

```

230     x = np.linspace(0, 90, len(mouse_1_day1_freezing))
231     freezing_data = mouse_1_day1_freezing
232     else:
233         x = np.linspace(0, 10, len(mouse_1_day6_freezing)) if day ==
234         6 else np.linspace(0, 10, len(mouse_1_day7_freezing))
235         freezing_data = mouse_1_day6_freezing if day == 6 else
236         mouse_1_day7_freezing
237
238     # Plot the freezing data on the bottom row axes
239     axes[0, i].plot(x, freezing_data, label="Freezing", color="r",
240                    alpha=0.6)
241     axes[0, i].set_title(f"Day {day} - Context A", fontsize=15)
242     axes[0, i].set_ylabel("Freezing (%)", fontsize=15)
243     axes[0, i].set_ylim(0, 100)
244     axes[0, i].legend()
245
246     if i != 0:
247         axes[0, i].set_yticklabels([])
248         axes[0, i].set_ylabel("")
249
250     if i == 1:
251         axes[0, i].set_title(f"Day {day} - Context B", fontsize=15)
252
253     if i == 2:
254         axes[0, i].set_title(f"Day {day} - Context B", fontsize=15)
255
256 plt.tight_layout()
257 plt.show()

```

D.2 TD-Momentum Model - Fitting SEFL Data

```

1 import numpy as np
2 from sklearn.preprocessing import MinMaxScaler
3 from scipy.optimize import minimize
4 from scipy.optimize import differential_evolution
5 import pandas as pd
6 import scipy.stats
7 import matplotlib.pyplot as plt
8
9 import scipy.io as sio
10 mat_contents = sio.loadmat('Supp_Mat')

```

```

11
12 # Copying versions of TD for use of TD Mom model and to stored
    results:
13
14 TD_Mom_stressed_results = pd.DataFrame.copy(TD_stressed_results)
15 TD_Mom_unstressed_results = pd.DataFrame.copy(TD_unstressed_results)
16
17 TD_Mom_stressed_results[TD_Mom_stressed_results.columns[1:]] = None
18 TD_Mom_unstressed_results[TD_Mom_stressed_results.columns[1:]] =
    None
19
20
21 ### Defining functions required ###
22 def combined_rescale(arr1, arr2, new_min, new_max):
23     combined_min = min(np.min(arr1), np.min(arr2))
24     combined_max = max(np.max(arr1), np.max(arr2))
25     rescaled_arr1 = (arr1 - combined_min) * (new_max - new_min) / (
    combined_max - combined_min + 1e-10) + new_min
26     rescaled_arr2 = (arr2 - combined_min) * (new_max - new_min) / (
    combined_max - combined_min + 1e-10) + new_min
27     return rescaled_arr1, rescaled_arr2
28
29 def TD_momentum_model_init(init, contexts, alpha, gamma1, gamma2, f)
    :
30     '''
31     Evaluates TD Momentum model threat at each timestep in each
    context only, no changing contexts, assumes the same
32     agent in one contex for the whole time, creates threat in other
    context as momentum term only.
33
34     :param init: Provides initialisation points for threat readings
    in both contexts (days 6 and 7)
35     :param contexts: Sequences of unconditioned stimuli across all
    contexts. e.g. sefl(sims = 2)[1] for 2nd mouse sim
36     :param u: Key of chosen context to compute TD momentum model for
    . (day1, day6, day7)
37     :param alpha: Learning rate
38     :param gamma1: Decay rate for threat
39     :param gamma2: Decay rate for momentum across all contexts
40     :param f: Scaling Constant for momentum term
41
42     :return: TD Momentum model threat estimation levels over all

```

```

timesteps
'''
43
44 # Initialise momentum array
45 m = np.zeros_like(contexts[:,0])
46
47 T_A = np.zeros_like(contexts[:,0])
48 T_B = np.zeros_like(contexts[:,1])
49
50 # Set initialisation
51 T_A[0] = init[0]
52 T_B[0] = init[1]
53
54 for t in range(1, len(contexts[:,0])):
55     # Set PE to 0 at every time step before computing PE for
current step
56     PE = 0
57     PE += alpha * (contexts[:, 0][t] - T_A[t - 1])
58     PE += alpha * (contexts[:, 1][t] - T_B[t - 1])
59
60     m[t] = m[t-1] + gamma2 * PE
61     m[t] = np.clip(m[t], 0, 1)
62
63     T_A[t] = T_A[t - 1] + alpha * (contexts[:, 0][t] - gamma1 *
T_A[t - 1]) + f * m[t]
64     T_B[t] = T_B[t - 1] + alpha * (contexts[:, 1][t] - gamma1 *
T_B[t - 1]) + f * m[t]
65
66     T_A[t] = np.clip(T_A[t], 0, 1)
67     T_B[t] = np.clip(T_B[t], 0, 1)
68
69     T_A, T_B = combined_rescale(T_A, T_B, 0.1, 0.9)
70
71     return T_A, T_B
72
73 def NLL_td_mom(params, init, stimuli, threats, u):
74     alpha, gamma1, gamma2, f = params
75     log_likelihood = 0.0
76
77     T_A, T_B = TD_momentum_model_init(init = init, contexts =
stimuli, alpha = alpha, gamma1 = gamma1, gamma2 = gamma2, f = f)
78
79     eps = 1e-10

```

```

80     T_A = np.clip(T_A, 0+eps, 1-eps)
81     T_B = np.clip(T_B, 0+eps, 1-eps)
82
83     if u == 0:
84         Tu = T_A
85
86     if u == 1:
87         Tu = T_B
88
89     for t in range(len(stimuli[:,0])):
90         log_likelihood += threats[t, u] * np.log(Tu[t]) + (1 -
91         threats[t, u]) * np.log(1 - Tu[t])
92
93     nll = -log_likelihood
94
95     return nll
96
97 def smoothing(raw_time_series, window_size):
98     smoothed_time_series = []
99     half_window = window_size // 2
100
101     # Pad the time series
102     padded_time_series = np.pad(raw_time_series, (half_window,
103     half_window), mode='edge')
104
105     for i in range(len(raw_time_series)):
106         window_values = padded_time_series[i : i + window_size]
107         smoothed_value = np.mean(window_values) # Average of window
108         smoothed_time_series.append(smoothed_value)
109
110     return np.array(smoothed_time_series)
111
112
113
114
115 ### FITTING CONTEXT A DAY 1 ### STRESSED
116
117 initial = [0,0]
118 initial_vals = np.array((0.5, 0.95, 0.4, 1.5))
119 bounds = [(0.05, 1), (0.9, 1), (0, 0.8), (0, 3)]

```

```

120
121 # Minimization for MICE
122 for m in range(0, len(stressed_mice_details)):
123     # Compute minimization for each participant
124     B_shocks = np.zeros_like(day1_stressed_attacks[m:m + 1].reshape
125                               (-1))
126     contexts_stacked = np.column_stack((day1_stressed_attacks[m:m +
127                               1].reshape(-1), B_shocks))
128
129     result_A = differential_evolution(NLL_td_mom, x0=initial_vals,
130                                     bounds=bounds,
131                                     args=(initial, contexts_stacked,
132                                           day1_stressed_freezing[m:m + 1].reshape((162000, 1)), 0),
133                                     polish=True)
134
135     print(f'A: n_iter: {result_A.nit} - success: {result_A.success}
136           - nll {result_A.fun}')
137
138     # Store in results dataframe:
139     TD_Mom_stressed_results["Day 1 Context A NLL"][m] = result_A.fun
140     alpha, gamma1, gamma2, f = result_A.x
141     TD_Mom_stressed_results["Day 1 Context A Params"][m] = [alpha,
142                                                             gamma1, gamma2, f]
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

```

```

152     a, g1, g2, f = TD_Mom_stressed_results["Day 1 Context A Params"
153     ][m]
154     day1_stressed_freezing_B = TD_momentum_model_init(init = initial
155     , contexts = contexts_stacked,
156
157
158
159
160
161
162
163
164
165
166
167     # Adjust NLL for day 6 - remove u param
168     def NLL_td_mom(params, init, stimuli, threats):
169         alpha, gamma1, gamma2, f = params
170         log_likelihood = 0.0
171
172         T_A, T_B = TD_momentum_model_init(init = init, contexts =
173         stimuli, alpha = alpha, gamma1 = gamma1, gamma2 = gamma2, f = f)
174
175         eps = 1e-10
176         T_A = np.clip(T_A, 0+eps, 1-eps)
177         T_B = np.clip(T_B, 0+eps, 1-eps)
178
179         Tu = T_B
180
181         for t in range(len(stimuli[:,0])):
182             log_likelihood += threats[t, 0] * np.log(Tu[t]) + (1 -
183             threats[t, 0]) * np.log(1 - Tu[t])
184
185         nll = -log_likelihood
186
187     return nll

```



```

187
188
189
190
191
192 ### FITTING CONTEXT B DAY 6 ### STRESSED
193
194 initial_vals = np.array((0.5, 0.95, 0.4, 1.5))
195 bounds = [(0.05, 1), (0.9, 1), (0, 0.8), (0, 3)]
196
197 # Minimization for MICE
198 for m in range(0, len(stressed_mice_details)):
199     # Setting up initial starting point for threat
200     a_prev, g1_prev, g2_prev, f_prev = TD_Mom_stressed_results["Day
201     1 Context A Params"][m]
202     initial_prev = [0,0]
203     B_shocks_prev = np.zeros_like(day1_stressed_attacks[m:m + 1].
204     reshape(-1))
205     contexts_stacked_prev = np.column_stack((day1_stressed_attacks[m
206     :m + 1].reshape(-1), B_shocks_prev))
207
208     initial_B = TD_momentum_model_init(init = initial_prev, contexts
209     = contexts_stacked_prev,
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999

```

```

218         polish=True)
219
220     print(f'A: n_iter: {result_B.nit} - success: {result_B.success}
221     - nll {result_B.fun}')
222
223     # Store in results dataframe:
224     TD_Mom_stressed_results["Day 6 Context B NLL"][m] = result_B.fun
225     alpha, gamma1, gamma2, f = result_B.x
226     TD_Mom_stressed_results["Day 6 Context B Params"][m] = [alpha,
227     gamma1, gamma2, f]
228
229     # Back to original NLL for day 6 context A
230     def NLL_td_mom(params, init, stimuli, threats, u):
231         alpha, gamma1, gamma2, f = params
232         log_likelihood = 0.0
233
234         T_A, T_B = TD_momentum_model_init(init = init, contexts =
235         stimuli, alpha = alpha, gamma1 = gamma1, gamma2 = gamma2, f = f)
236
237         eps = 1e-10
238         T_A = np.clip(T_A, 0+eps, 1-eps)
239         T_B = np.clip(T_B, 0+eps, 1-eps)
240
241         if u == 0:
242             Tu = T_A
243
244         if u == 1:
245             Tu = T_B
246
247         for t in range(len(stimuli[:,0])):
248             log_likelihood += threats[t, u] * np.log(Tu[t]) + (1 -
249             threats[t, u]) * np.log(1 - Tu[t])
250
251         nll = -log_likelihood
252
253         return nll
254
255     ### FITTING CONTEXT A DAY 6 ### STRESSED
256
257     # Minimization for MICE
258     for m in range(0, len(stressed_mice_details)):

```

```

256     # Setting up initial starting point for threat
257     a_prev, g1_prev, g2_prev, f_prev = TD_Mom_stressed_results["Day
258     1 Context A Params"][m]
259     initial_prev = [0,0]
260     B_shocks_prev = np.zeros_like(day1_stressed_attacks[m:m + 1].
261     reshape(-1))
262     contexts_stacked_prev = np.column_stack((day1_stressed_attacks[m
263     :m + 1].reshape(-1), B_shocks_prev))
264
265     initial_A = TD_momentum_model_init(init = initial_prev, contexts
266     = contexts_stacked_prev,
267
268                                     alpha=a,
269
270     gamma1=g1, gamma2=g2, f=f)[0][-1]
271
272     initial_B = TD_momentum_model_init(init = initial_prev, contexts
273     = contexts_stacked_prev,
274
275                                     alpha=a,
276
277     gamma1=g1, gamma2=g2, f=f)[1][-1]
278
279     # Compute TD Mom threat fit from parameters for this day,
280     compute NLL
281     A_shocks = np.zeros_like(day6_stressed_attacks_B[m:m + 1].
282     reshape(-1))
283     contexts_stacked = np.column_stack((A_shocks,
284     day6_stressed_attacks_B[m:m + 1].reshape(-1)))
285     initial = [initial_A, initial_B]
286
287     a, g1, g2, f = TD_Mom_stressed_results["Day 6 Context B Params"
288     ][m]
289     day6_stressed_freezing_A = TD_momentum_model_init(init = initial
290     , contexts = contexts_stacked,
291
292                                     alpha=a,
293
294     gamma1=g1, gamma2=g2, f=f)[0]
295
296     params = [a, g1, g2, f]
297     threats_stacked = np.column_stack((day6_stressed_freezing_A,
298     day6_stressed_freezing[m]))
299     nll = NLL_td_mom(params=params, init=initial, stimuli=
300     contexts_stacked, threats=threats_stacked, u=0)
301     # Store in results dataframe:
302     TD_Mom_stressed_results["Day 6 Context A NLL"][m] = nll
303
304

```

```

283
284
285
286
287 # Adjust NLL for day 7 - remove u param
288 def NLL_td_mom(params, init, stimuli, threats):
289     alpha, gammal, gamma2, f = params
290     log_likelihood = 0.0
291
292     T_A, T_B = TD_momentum_model_init(init = init, contexts =
stimuli, alpha = alpha, gammal = gammal, gamma2 = gamma2, f = f)
293
294     eps = 1e-10
295     T_A = np.clip(T_A, 0+eps, 1-eps)
296     T_B = np.clip(T_B, 0+eps, 1-eps)
297
298
299     Tu = T_B
300
301     for t in range(len(stimuli[:,0])):
302         log_likelihood += threats[t, 0] * np.log(Tu[t]) + (1 -
threats[t, 0]) * np.log(1 - Tu[t])
303
304     nll = -log_likelihood
305
306     return nll
307
308
309 ### FITTING CONTEXT B DAY 7 ### STRESSED
310
311
312 initial_vals = np.array((0.5, 0.95, 0.4, 1.5))
313 bounds = [(0.05, 1), (0.9, 1), (0, 0.8), (0, 3)]
314
315 # Minimization for MICE
316 for m in range(0, len(stressed_mice_details)):
317     # Setting up initial starting point for threat
318     a_prev_prev, g1_prev_prev, g2_prev_prev, f_prev_prev =
TD_Mom_stressed_results["Day 1 Context A Params"][m]
319     initial_prev_prev = [0,0]
320     B_shocks_prev_prev = np.zeros_like(day1_stressed_attacks[m:m +
1].reshape(-1))

```

```

321     contexts_stacked_prev_prev = np.column_stack((
day1_stressed_attacks[m:m + 1].reshape(-1), B_shocks_prev_prev))
322
323     initial_prev_A = TD_momentum_model_init(init = initial_prev_prev
, contexts = contexts_stacked_prev_prev,
324
alpha=
a_prev_prev, gamma1=g1_prev_prev,
325
gamma2=g2_prev_prev, f=
f_prev_prev)[0][-1]
326
327     initial_prev_B = TD_momentum_model_init(init = initial_prev_prev
, contexts = contexts_stacked_prev_prev,
328
alpha=
a_prev_prev, gamma1=g1_prev_prev,
329
gamma2=g2_prev_prev, f=
f_prev_prev)[1][-1]
330
331
332     initial_prev = [initial_prev_A, initial_prev_B]
333     a_prev, g1_prev, g2_prev, f_prev = TD_Mom_stressed_results["Day
6 Context B Params"][m]
334     A_shocks_prev = np.zeros_like(day6_stressed_attacks_B[m:m + 1].
reshape(-1))
335     contexts_stacked_prev = np.column_stack((A_shocks_prev,
day6_stressed_attacks_B[m:m + 1].reshape(-1)))
336
337     initial_A = TD_momentum_model_init(init=initial_prev, contexts=
contexts_stacked_prev,
338
alpha=a_prev, gamma1=
g1_prev,
339
gamma2=g2_prev, f=f_prev)
[0][-1]
340
341     initial_B = TD_momentum_model_init(init=initial_prev, contexts=
contexts_stacked_prev,
342
alpha=a_prev, gamma1=
g1_prev,
343
gamma2=g2_prev, f=f_prev)
[1][-1]
344
345     # Compute minimization for each participant
346     initial = [initial_A, initial_B]

```

```

347     A_shocks = np.zeros_like(day7_stressed_attacks_A[m:m + 1].
reshape(-1))    # doesn't matter if A or B, both filled with
zeros
348     contexts_stacked = np.column_stack((A_shocks,
day7_stressed_attacks_B[m:m + 1].reshape(-1)))
349
350     result_B = differential_evolution(NLL_td_mom, x0=initial_vals,
bounds=bounds,
351                                     args=(initial, contexts_stacked,
day7_stressed_freezing[m:m + 1].reshape((18000, 1))),
352                                     polish=True)
353
354     print(f'B: n_iter: {result_B.nit} - success: {result_B.success}
- nll {result_B.fun}')
355
356     # Store in results dataframe:
357     TD_Mom_stressed_results["Day 7 Context B NLL"][m] = result_B.fun
358     alpha, gamma1, gamma2, f = result_B.x
359     TD_Mom_stressed_results["Day 7 Context A Params"][m] = [alpha,
gamma1, gamma2, f]
360
361
362
363
364
365 # Back to original NLL for day 7 context A
366 def NLL_td_mom(params, init, stimuli, threats, u):
367     alpha, gamma1, gamma2, f = params
368     log_likelihood = 0.0
369
370     T_A, T_B = TD_momentum_model_init(init = init, contexts =
stimuli, alpha = alpha, gamma1 = gamma1, gamma2 = gamma2, f = f)
371
372     eps = 1e-10
373     T_A = np.clip(T_A, 0+eps, 1-eps)
374     T_B = np.clip(T_B, 0+eps, 1-eps)
375
376     if u == 0:
377         Tu = T_A
378
379     if u == 1:
380         Tu = T_B

```

```

381
382     for t in range(len(stimuli[:,0])):
383         log_likelihood += threats[t, u] * np.log(Tu[t]) + (1 -
384             threats[t, u]) * np.log(1 - Tu[t])
385
386     nll = -log_likelihood
387
388     return nll
389
390 ### FITTING CONTEXT A DAY 7 ### STRESSED
391
392 for m in range(0, len(stressed_mice_details)):
393     # Setting up initial starting point for threat
394     a_prev_prev, g1_prev_prev, g2_prev_prev, f_prev_prev =
395     TD_Mom_stressed_results["Day 1 Context A Params"][m]
396     initial_prev_prev = [0, 0]
397     B_shocks_prev_prev = np.zeros_like(day1_stressed_attacks[m:m +
398         1].reshape(-1))
399     contexts_stacked_prev_prev = np.column_stack((
400         day1_stressed_attacks[m:m + 1].reshape(-1), B_shocks_prev_prev))
401
402     initial_prev_A = TD_momentum_model_init(init=initial_prev_prev,
403         contexts=contexts_stacked_prev_prev,
404
405         alpha=a_prev_prev, gamma1=
406         g1_prev_prev,
407
408         gamma2=g2_prev_prev, f=
409         f_prev_prev)[0][-1]
410
411     initial_prev_B = TD_momentum_model_init(init=initial_prev_prev,
412         contexts=contexts_stacked_prev_prev,
413
414         alpha=a_prev_prev, gamma1=
415         g1_prev_prev,
416
417         gamma2=g2_prev_prev, f=
418         f_prev_prev)[1][-1]
419
420     a_prev, g1_prev, g2_prev, f_prev = TD_Mom_stressed_results["Day
421     6 Context B Params"][m]
422     A_shocks_prev = np.zeros_like(day6_stressed_attacks_B[m:m + 1].
423         reshape(-1))
424     contexts_stacked_prev = np.column_stack((A_shocks_prev,
425         day6_stressed_attacks_B[m:m + 1].reshape(-1)))
426     initial_prev = [initial_prev_A, initial_prev_B]
427
428

```

```

410     initial_A = TD_momentum_model_init(init=initial_prev, contexts=
contexts_stacked_prev,
411                                     alpha=a_prev, gammad1=g1_prev,
412                                     gamma2=g2_prev, f=f_prev)
[0][-1]
413
414     initial_B = TD_momentum_model_init(init=initial_prev, contexts=
contexts_stacked_prev,
415                                     alpha=a_prev, gammad1=g1_prev,
416                                     gamma2=g2_prev, f=f_prev)
[1][-1]
417
418     # Compute minimization for each participant
419     A_shocks = np.zeros_like(day7_stressed_attacks_A[m:m + 1].
reshape(-1))
420     contexts_stacked = np.column_stack((A_shocks,
day7_stressed_attacks_B[m:m + 1].reshape(-1)))
421     a, g1, g2, f = TD_Mom_stressed_results["Day 7 Context A Params"
][m]
422
423     initial = [initial_A, initial_B]
424
425     day7_stressed_freezing_A = TD_momentum_model_init(init=initial,
contexts=contexts_stacked, alpha=a, gammad1=g1,
426                                                         gamma2=g2, f=f
) [0]
427
428     params = [a, g1, g2, f]
429     threats_stacked = np.column_stack((day7_stressed_freezing_A,
day7_stressed_freezing[m]))
430     nll = NLL_td_mom(params=params, init=initial, stimuli=
contexts_stacked, threats=threats_stacked, u=0)
431     # Store in results dataframe:
432     TD_Mom_stressed_results["Day 7 Context A NLL"][m] = nll
433
434
435
436
437     ### NOW FOR UNSTRESSED ###
438
439
440

```



```

441 ### FITTING CONTEXT A DAY 1 ### UNSTRESSED
442
443 initial = [0,0]
444 initial_vals = np.array((0.5, 0.95, 0.4, 1.5))
445 bounds = [(0.05, 1), (0.9, 1), (0, 0.8), (0, 3)]
446
447 # Minimization for MICE
448 for m in range(0, len(unstressed_mice_details)):
449     # Compute minimization for each participant
450     B_shocks = np.zeros_like(day1_unstressed_attacks[m:m + 1].
451     reshape(-1))
452     contexts_stacked = np.column_stack((day1_unstressed_attacks[m:m
453     + 1].reshape(-1), B_shocks))
454
455     result_A = differential_evolution(NLL_td_mom, x0=initial_vals,
456     bounds=bounds,
457     args=(initial, contexts_stacked,
458     day1_unstressed_freezing[m:m + 1].reshape((162000, 1)), 0),
459     polish=True)
460
461     print(f'A: n_iter: {result_A.nit} - success: {result_A.success}
462     - nll {result_A.fun}')
463
464     # Store in results dataframe:
465     TD_Mom_unstressed_results["Day 1 Context A NLL"][m] = result_A.
466     fun
467     alpha, gamma1, gamma2, f = result_A.x
468     TD_Mom_unstressed_results["Day 1 Context A Params"][m] = [alpha,
469     gamma1, gamma2, f]
470
471 ### FITTING CONTEXT B DAY 1 ### UNSTRESSED
472 day1_unstressed_attacks_B = np.zeros_like(day1_unstressed_attacks)
473     # (15, 162000) use as stimuli input
474
475 initial = [0,0]
476
477 # Minimization for MICE
478 for m in range(0, len(unstressed_mice_details)):
479     # Compute minimization for each participant
480     B_shocks = np.zeros_like(day1_unstressed_attacks[m:m + 1].

```

```

reshape(-1))
475 contexts_stacked = np.column_stack((day1_unstressed_attacks[m:m
+ 1].reshape(-1), B_shocks))
476
477 # Compute TD threat fit from parameters for this day, compute
NLL
478 a, g1, g2, f = TD_Mom_unstressed_results["Day 1 Context A Params
"] [m]
479 day1_unstressed_freezing_B = TD_momentum_model_init(init =
initial, contexts = contexts_stacked,
480
alpha=a,
gamma1=g1, gamma2=g2, f=f) [1]
481
482 params = [a, g1, g2, f]
483 threats_stacked = np.column_stack((day1_unstressed_freezing[m],
day1_unstressed_freezing_B))
484 nll = NLL_td_mom(params=params, init=initial, stimuli=
contexts_stacked, threats=threats_stacked, u=1)
485 # Store in results dataframe:
486 TD_Mom_unstressed_results["Day 1 Context B NLL"] [m] = nll
487
488
489
490
491
492
493 # Adjust NLL for day 6 - remove u param
494 def NLL_td_mom(params, init, stimuli, threats):
495     alpha, gamma1, gamma2, f = params
496     log_likelihood = 0.0
497
498     T_A, T_B = TD_momentum_model_init(init = init, contexts =
stimuli, alpha = alpha, gamma1 = gamma1, gamma2 = gamma2, f = f)
499
500     eps = 1e-10
501     T_A = np.clip(T_A, 0+eps, 1-eps)
502     T_B = np.clip(T_B, 0+eps, 1-eps)
503
504
505     Tu = T_B
506
507     for t in range(len(stimuli[:,0])):

```

```

508     log_likelihood += threats[t, 0] * np.log(Tu[t]) + (1 -
threats[t, 0]) * np.log(1 - Tu[t])
509
510     nll = -log_likelihood
511
512     return nll
513
514
515
516
517
518 ### FITTING CONTEXT B DAY 6 ### UNSTRESSED
519
520 initial_vals = np.array((0.5, 0.95, 0.4, 1.5))
521 bounds = [(0.05, 1), (0.9, 1), (0, 0.8), (0, 3)]
522
523 # Minimization for MICE
524 for m in range(0, len(unstressed_mice_details)):
525     # Setting up initial starting point for threat
526     a_prev, g1_prev, g2_prev, f_prev = TD_Mom_unstressed_results["
Day 1 Context A Params"][m]
527     initial_prev = [0,0]
528     B_shocks_prev = np.zeros_like(day1_unstressed_attacks[m:m + 1].
reshape(-1))
529     contexts_stacked_prev = np.column_stack((day1_unstressed_attacks
[m:m + 1].reshape(-1), B_shocks_prev))
530
531     initial_B = TD_momentum_model_init(init = initial_prev, contexts
= contexts_stacked_prev,
532                                     alpha=a,
gamma1=g1, gamma2=g2, f=f)[1][-1]
533
534     initial_A = TD_momentum_model_init(init = initial_prev, contexts
= contexts_stacked_prev,
535                                     alpha=a,
gamma1=g1, gamma2=g2, f=f)[0][-1]
536
537     # Compute minimization for each participant
538     A_shocks = np.zeros_like(day6_unstressed_attacks_B[m:m + 1].
reshape(-1))
539     contexts_stacked = np.column_stack((A_shocks,
day6_unstressed_attacks_B[m:m + 1].reshape(-1)))

```

```

540     initial = [initial_A, initial_B]
541
542     result_B = differential_evolution(NLL_td_mom, x0=initial_vals,
543     bounds=bounds,
544     args=(initial, contexts_stacked,
545     day6_unstressed_freezing[m:m + 1].reshape((18000, 1))),
546     polish=True)
547
548     print(f'A: n_iter: {result_B.nit} - success: {result_B.success}
549     - nll {result_B.fun}')
550
551     # Store in results dataframe:
552     TD_Mom_unstressed_results["Day 6 Context B NLL"][m] = result_B.
553     fun
554     alpha, gamma1, gamma2, f = result_B.x
555     TD_Mom_unstressed_results["Day 6 Context B Params"][m] = [alpha,
556     gamma1, gamma2, f]
557
558
559     # Back to original NLL for day 6 context A
560     def NLL_td_mom(params, init, stimuli, threats, u):
561         alpha, gamma1, gamma2, f = params
562         log_likelihood = 0.0
563
564         T_A, T_B = TD_momentum_model_init(init = init, contexts =
565         stimuli, alpha = alpha, gamma1 = gamma1, gamma2 = gamma2, f = f)
566
567         eps = 1e-10
568         T_A = np.clip(T_A, 0+eps, 1-eps)
569         T_B = np.clip(T_B, 0+eps, 1-eps)
570
571         if u == 0:
572             Tu = T_A
573
574         if u == 1:
575             Tu = T_B
576
577         for t in range(len(stimuli[:,0])):
578             log_likelihood += threats[t, u] * np.log(Tu[t]) + (1 -
579             threats[t, u]) * np.log(1 - Tu[t])
580
581         nll = -log_likelihood

```

```

575
576     return nll
577
578     ### FITTING CONTEXT A DAY 6 ### UNSTRESSED
579
580     # Minimization for MICE
581     for m in range(0, len(unstressed_mice_details)):
582         # Setting up initial starting point for threat
583         a_prev, g1_prev, g2_prev, f_prev = TD_Mom_unstressed_results["
Day 1 Context A Params"][m]
584         initial_prev = [0,0]
585         B_shocks_prev = np.zeros_like(day1_unstressed_attacks[m:m + 1].
reshape(-1))
586         contexts_stacked_prev = np.column_stack((day1_unstressed_attacks
[m:m + 1].reshape(-1), B_shocks_prev))
587
588         initial_A = TD_momentum_model_init(init = initial_prev, contexts
= contexts_stacked_prev,
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
                    alpha=a,
                    gamma1=g1, gamma2=g2, f=f)[0][-1]
                    initial_B = TD_momentum_model_init(init = initial_prev, contexts
= contexts_stacked_prev,
                    alpha=a,
                    gamma1=g1, gamma2=g2, f=f)[1][-1]
                    # Compute TD Mom threat fit from parameters for this day,
                    compute NLL
                    A_shocks = np.zeros_like(day6_unstressed_attacks_B[m:m + 1].
reshape(-1))
                    contexts_stacked = np.column_stack((A_shocks,
day6_unstressed_attacks_B[m:m + 1].reshape(-1)))
                    initial = [initial_A, initial_B]
                    a, g1, g2, f = TD_Mom_unstressed_results["Day 6 Context B Params
"][m]
                    day6_unstressed_freezing_A = TD_momentum_model_init(init =
initial, contexts = contexts_stacked,
                    alpha=a,
                    gamma1=g1, gamma2=g2, f=f)[0]
                    params = [a, g1, g2, f]

```

```

604     threats_stacked = np.column_stack((day6_unstressed_freezing_A,
605     day6_unstressed_freezing[m]))
606     nll = NLL_td_mom(params=params, init=initial, stimuli=
607     contexts_stacked, threats=threats_stacked, u=0)
608     # Store in results dataframe:
609     TD_Mom_unstressed_results["Day 6 Context A NLL"][m] = nll
610
611
612
613
614
615
616
617
618 # Adjust NLL for day 7 - remove u param
619 def NLL_td_mom(params, init, stimuli, threats):
620     alpha, gamma1, gamma2, f = params
621     log_likelihood = 0.0
622
623     T_A, T_B = TD_momentum_model_init(init = init, contexts =
624     stimuli, alpha = alpha, gamma1 = gamma1, gamma2 = gamma2, f = f)
625
626     eps = 1e-10
627     T_A = np.clip(T_A, 0+eps, 1-eps)
628     T_B = np.clip(T_B, 0+eps, 1-eps)
629
630     Tu = T_B
631
632     for t in range(len(stimuli[:,0])):
633         log_likelihood += threats[t, 0] * np.log(Tu[t]) + (1 -
634         threats[t, 0]) * np.log(1 - Tu[t])
635
636     nll = -log_likelihood
637
638     return nll
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

```

```

642 initial_vals = np.array((0.5, 0.95, 0.4, 1.5))
643 bounds = [(0.05, 1), (0.9, 1), (0, 0.8), (0, 3)]
644
645 # Minimization for MICE
646 for m in range(0, len(unstressed_mice_details)):
647     # Setting up initial starting point for threat
648     a_prev_prev, g1_prev_prev, g2_prev_prev, f_prev_prev =
TD_Mom_unstressed_results["Day 1 Context A Params"][m]
649     initial_prev_prev = [0,0]
650     B_shocks_prev_prev = np.zeros_like(day1_unstressed_attacks[m:m +
1]).reshape(-1))
651     contexts_stacked_prev_prev = np.column_stack((
day1_unstressed_attacks[m:m + 1].reshape(-1), B_shocks_prev_prev)
)
652
653     initial_prev_A = TD_momentum_model_init(init = initial_prev_prev
, contexts = contexts_stacked_prev_prev,
654
alpha=
a_prev_prev, gammal=g1_prev_prev,
655
gamma2=g2_prev_prev, f=
f_prev_prev)[0][-1]
656
657     initial_prev_B = TD_momentum_model_init(init = initial_prev_prev
, contexts = contexts_stacked_prev_prev,
658
alpha=
a_prev_prev, gammal=g1_prev_prev,
659
gamma2=g2_prev_prev, f=
f_prev_prev)[1][-1]
660
661
662     initial_prev = [initial_prev_A, initial_prev_B]
663     a_prev, g1_prev, g2_prev, f_prev = TD_Mom_unstressed_results["
Day 6 Context B Params"][m]
664     A_shocks_prev = np.zeros_like(day6_unstressed_attacks_B[m:m +
1]).reshape(-1))
665     contexts_stacked_prev = np.column_stack((A_shocks_prev,
day6_unstressed_attacks_B[m:m + 1].reshape(-1)))
666
667     initial_A = TD_momentum_model_init(init=initial_prev, contexts=
contexts_stacked_prev,
668
alpha=a_prev, gammal=
g1_prev,

```

```

669                                     gamma2=g2_prev, f=f_prev)
670 [0][-1]
671     initial_B = TD_momentum_model_init(init=initial_prev, contexts=
contexts_stacked_prev,
672                                     alpha=a_prev, gamma1=
gl_prev,
673                                     gamma2=g2_prev, f=f_prev)
674 [1][-1]
675     # Compute minimization for each participant
676     initial = [initial_A, initial_B]
677     A_shocks = np.zeros_like(day7_unstressed_attacks_A[m:m + 1].
reshape(-1))    # doesn't matter if A or B, both filled with
zeros
678     contexts_stacked = np.column_stack((A_shocks,
day7_unstressed_attacks_B[m:m + 1].reshape(-1)))
679
680     result_B = differential_evolution(NLL_td_mom, x0=initial_vals,
bounds=bounds,
681                                     args=(initial, contexts_stacked,
day7_unstressed_freezing[m:m + 1].reshape((18000, 1))),
682                                     polish=True)
683
684     print(f'B: n_iter: {result_B.nit} - success: {result_B.success}
- nll {result_B.fun}')
685
686     # Store in results dataframe:
687     TD_Mom_unstressed_results["Day 7 Context B NLL"][m] = result_B.
fun
688     alpha, gamma1, gamma2, f = result_B.x
689     TD_Mom_unstressed_results["Day 7 Context A Params"][m] = [alpha,
gamma1, gamma2, f]
690
691
692
693
694
695 # Back to original NLL for day 7 context A
696 def NLL_td_mom(params, init, stimuli, threats, u):
697     alpha, gamma1, gamma2, f = params
698     log_likelihood = 0.0

```



```

699
700     T_A, T_B = TD_momentum_model_init(init = init, contexts =
stimuli, alpha = alpha, gammal = gammal, gamma2 = gamma2, f = f)
701
702     eps = 1e-10
703     T_A = np.clip(T_A, 0+eps, 1-eps)
704     T_B = np.clip(T_B, 0+eps, 1-eps)
705
706     if u == 0:
707         Tu = T_A
708
709     if u == 1:
710         Tu = T_B
711
712     for t in range(len(stimuli[:,0])):
713         log_likelihood += threats[t, u] * np.log(Tu[t]) + (1 -
threats[t, u]) * np.log(1 - Tu[t])
714
715     nll = -log_likelihood
716
717     return nll
718 ### FITTING CONTEXT A DAY 7 ### UNSTRESSED
719
720 for m in range(0, len(unstressed_mice_details)):
721     # Setting up initial starting point for threat
722     a_prev_prev, g1_prev_prev, g2_prev_prev, f_prev_prev =
TD_Mom_unstressed_results["Day 1 Context A Params"][m]
723     initial_prev_prev = [0, 0]
724     B_shocks_prev_prev = np.zeros_like(day1_unstressed_attacks[m:m +
1]).reshape(-1)
725     contexts_stacked_prev_prev = np.column_stack((
day1_unstressed_attacks[m:m + 1].reshape(-1), B_shocks_prev_prev)
)
726
727     initial_prev_A = TD_momentum_model_init(init=initial_prev_prev,
contexts=contexts_stacked_prev_prev,
728                                             alpha=a_prev_prev, gammal=
g1_prev_prev,
729                                             gamma2=g2_prev_prev, f=
f_prev_prev)[0][-1]
730
731     initial_prev_B = TD_momentum_model_init(init=initial_prev_prev,

```

```

contexts=contexts_stacked_prev_prev,
732                                     alpha=a_prev_prev, gamma1=
g1_prev_prev,
733                                     gamma2=g2_prev_prev, f=
f_prev_prev)[1][-1]
734
735 a_prev, g1_prev, g2_prev, f_prev = TD_Mom_unstressed_results["
Day 6 Context B Params"][m]
736 A_shocks_prev = np.zeros_like(day6_unstressed_attacks_B[m:m +
1]).reshape(-1)
737 contexts_stacked_prev = np.column_stack((A_shocks_prev,
day6_unstressed_attacks_B[m:m + 1].reshape(-1)))
738 initial_prev = [initial_prev_A, initial_prev_B]
739
740 initial_A = TD_momentum_model_init(init=initial_prev, contexts=
contexts_stacked_prev,
741                                     alpha=a_prev, gamma1=g1_prev,
742                                     gamma2=g2_prev, f=f_prev)
[0][-1]
743
744 initial_B = TD_momentum_model_init(init=initial_prev, contexts=
contexts_stacked_prev,
745                                     alpha=a_prev, gamma1=g1_prev,
746                                     gamma2=g2_prev, f=f_prev)
[1][-1]
747
748 # Compute minimization for each participant
749 A_shocks = np.zeros_like(day7_unstressed_attacks_A[m:m + 1]).
reshape(-1)
750 contexts_stacked = np.column_stack((A_shocks,
day7_unstressed_attacks_B[m:m + 1].reshape(-1)))
751 a, g1, g2, f = TD_Mom_unstressed_results["Day 7 Context A Params
"][m]
752
753 initial = [initial_A, initial_B]
754
755 day7_stressed_freezing_A = TD_momentum_model_init(init=initial,
contexts=contexts_stacked, alpha=a, gamma1=g1,
756                                     gamma2=g2, f=f
) [0]
757
758 params = [a, g1, g2, f]

```

```

759     threats_stacked = np.column_stack((day7_unstressed_freezing_A,
760     day7_unstressed_freezing[m]))
761     nll = NLL_td_mom(params=params, init=initial, stimuli=
762     contexts_stacked, threats=threats_stacked, u=0)
761     # Store in results dataframe:
762     TD_Mom_unstressed_results["Day 7 Context A NLL"][m] = nll

```

D.3 TD-Momentum Model - Computing BIC Scores

```

1   ### BIC SCORES FOR TD Momentum MODEL ###
2
3   # BIC = 2 * NLL + p*log(n)
4
5   # where p = number of params, n = number of observations (162000
6     day1 and 18000 for day 6 and 7)
7
8   p = 4
9
10  # Stressed BIC
11  for m in range(len(TD_Mom_stressed_results)):
12
13      Day1_NLL = (TD_Mom_stressed_results["Day 1 Context A NLL"][m] +
14      TD_Mom_stressed_results["Day 1 Context B NLL"][m])
15
16      Day6_NLL = (TD_Mom_stressed_results["Day 6 Context A NLL"][m] +
17      TD_Mom_stressed_results["Day 6 Context B NLL"][m])
18
19      Day7_NLL = (TD_Mom_stressed_results["Day 7 Context A NLL"][m] +
20      TD_Mom_stressed_results["Day 7 Context B NLL"][m])
21
22      Day1_BIC = 2 * Day1_NLL + p * np.log(162000)
23
24      Day6_BIC = 2 * Day6_NLL + p * np.log(18000)
25
26      Day7_BIC = 2 * Day7_NLL + p * np.log(18000)
27
28      TD_Mom_stressed_results["BIC Score"][m] = Day1_BIC + Day6_BIC +
29      Day7_BIC
30
31  # Unstressed BIC
32  for m in range(len(TD_Mom_unstressed_results)):
33      Day1_NLL = (TD_Mom_unstressed_results["Day 1 Context A NLL"][m]
34      +

```

```

30         TD_Mom_unstressed_results["Day 1 Context B NLL"][m])
31
32     Day6_NLL = (TD_Mom_unstressed_results["Day 6 Context A NLL"][m]
33 +
34         TD_Mom_unstressed_results["Day 6 Context B NLL"][m])
35
36     Day7_NLL = (TD_Mom_unstressed_results["Day 7 Context A NLL"][m]
37 +
38         TD_Mom_unstressed_results["Day 7 Context B NLL"][m])
39
40     Day1_BIC = 2 * Day1_NLL + p * np.log(162000)
41
42     Day6_BIC = 2 * Day6_NLL + p * np.log(18000)
43
44     Day7_BIC = 2 * Day7_NLL + p * np.log(18000)
45
46     TD_Mom_unstressed_results["BIC Score"][m] = Day1_BIC + Day6_BIC
47 + Day7_BIC

```

D.4 TD-Momentum Model - Multivariate Sampled Parameter Recovery

```

1 import matplotlib.pyplot as plt
2 import numpy as np
3 from scipy.optimize import minimize
4 import scipy
5 import pandas as pd
6 from scipy.optimize import differential_evolution
7
8
9 # Model Param lists (taken from paper supplementary material)
10 # Loading matlab file from paper - creates dict "mat_contents" which
11     stores all variables
12 import scipy.io as sio
13 mat_contents = sio.loadmat('Supp_Mat')
14
15 def day_1(shocks):
16     '''
17     Evaluates day1 of SEFL experiment.

```

```

18     :return: Outputs sequence of 15 shocks randomly over 90 mins
        (5400 secs)
19     '''
20     min_shock_int = 4 * 1800
21     max_shock_int = 6 * 1800
22     # 90 mins = 5400 secs
23     day_1 = np.zeros(162000)
24
25     all_shock_intervals = []
26     while np.sum(day_1) < shocks:
27         # Generate random interval between 4 and 8 mins
28         shock_interval = np.random.randint(min_shock_int,
        max_shock_int)
29         # Add this to the list of all interval times
30         all_shock_intervals.append(shock_interval)
31         # Calculate the shock index by summing cumulative intervals
32         shock_index = np.sum(all_shock_intervals)
33         # Add shock to day_1 array
34         day_1[shock_index] = 1
35     return day_1
36
37 # Combined scaling function for context A and B
38 def combined_rescale(arr1, arr2, new_min, new_max):
39     combined_min = min(np.min(arr1), np.min(arr2))
40     combined_max = max(np.max(arr1), np.max(arr2))
41     rescaled_arr1 = (arr1 - combined_min) * (new_max - new_min) / (
        combined_max - combined_min + 1e-10) + new_min
42     rescaled_arr2 = (arr2 - combined_min) * (new_max - new_min) / (
        combined_max - combined_min + 1e-10) + new_min
43     return rescaled_arr1, rescaled_arr2
44
45 # Creating NLL function and performing parameter recovery for
        simulated data
46 def TD_momentum_model(contexts, alpha, gamma1, gamma2, f):
47     '''
48     Evaluates TD Momentum model threat at each timestep in each
        context only, no changing contexts, assumes the same
49     agent in one context for the whole time, creates threat in other
        context as momentum term only.
50
51     :param contexts: Sequences of unconditioned stimuli across all
        contexts. e.g. sefl(sims = 2)[1] for 2nd mouse sim

```

```

52     :param u: Key of chosen context to compute TD momentum model for
      . (day1, day6, day7)
53     :param alpha: Learning rate
54     :param gamma1: Decay rate for threat
55     :param gamma2: Decay rate for momentum across all contexts
56     :param f: Scaling Constant for momentum term
57
58     :return: TD Momentum model threat estimation levels over all
      timesteps
59     '''
60     # Initialise momentum array
61     m = np.zeros_like(contexts[:,0])
62
63     T_A = np.zeros_like(contexts[:,0])
64     T_B = np.zeros_like(contexts[:,1])
65
66
67     for t in range(1, len(contexts[:,0])):
68         # Set PE to 0 at every time step before computing PE for
      current step
69         PE = 0
70         PE += alpha * (contexts[:, 0][t] - T_A[t - 1])
71         PE += alpha * (contexts[:, 1][t] - T_B[t - 1])
72
73         m[t] = m[t-1] + gamma2 * PE
74         m[t] = np.clip(m[t], 0, 1)
75
76         T_A[t] = T_A[t - 1] + alpha * (contexts[:, 0][t] - gamma1 *
      T_A[t - 1]) + f * m[t]
77         T_B[t] = T_B[t - 1] + alpha * (contexts[:, 1][t] - gamma1 *
      T_B[t - 1]) + f * m[t]
78
79         T_A[t] = np.clip(T_A[t], 0, 1)
80         T_B[t] = np.clip(T_B[t], 0, 1)
81
82     #T_A, T_B = combined_rescale(T_A, T_B, 0.1, 0.9)
83
84     return T_A, T_B
85
86
87 def NLL_td_mom(params, stimuli, threats, u):
88     alpha, gamma1, gamma2, f = params

```

```
89     log_likelihood = 0.0
90
91     T_A, T_B = TD_momentum_model(contexts = stimuli, alpha = alpha,
92     gamma1 = gamma1, gamma2 = gamma2, f = f)
93
94     eps = 1e-10
95     T_A = np.clip(T_A, 0+eps, 1-eps)
96     T_B = np.clip(T_B, 0+eps, 1-eps)
97
98     if u == 0:
99         Tu = T_A
100
101     if u == 1:
102         Tu = T_B
103
104     for t in range(len(stimuli[:,0])):
105         log_likelihood += threats[t, u] * np.log(Tu[t]) + (1 -
106         threats[t, u]) * np.log(1 - Tu[t])
107
108     nll = -log_likelihood
109
110     return nll
111
112 ### PARAMETER RECOVERY FOR 100 MULTIVARIATE NORMAL SAMPLES PARAMETER
113     SETS ### CONTEXT A ###
114 bounds = [(0.05,1), (0.9,1), (0,0.8), (0,3)]
115
116 alpha_mean = 0.5
117 gamma1_mean = 0.95
118 gamma2_mean = 0.45
119 f_mean = 1.5
120 means = [alpha_mean, gamma1_mean, gamma2_mean, f_mean]
121
122 alpha_var = 0.05
123 gamma1_var = 0.001
124 gamma2_var = 0.05
125 f_var = 0.5
126
127 covs = np.zeros((4,4))
128 covs[0,0] = alpha_var
129 covs[1,1] = gamma1_var
```

```

128 covs[2,2] = gamma2_var
129 covs[3,3] = f_var
130
131 sample_params = np.random.multivariate_normal(mean=means, cov=covs,
        size=100)
132 sample_params[:,0] = np.clip(sample_params[:,0], 0.05, 0.9)
133 sample_params[:,1] = np.clip(sample_params[:,1], 0.9, 1)
134 sample_params[:,2] = np.clip(sample_params[:,2], 0, 0.8)
135 sample_params[:,3] = np.clip(sample_params[:,3], 0, 3)
136
137 '''
138 # Plot of param values
139 fig, ax = plt.subplots(4)
140 fig.suptitle("Sampled Parameter Values")
141 # add a big axes, hide frame
142 fig.add_subplot(111, frameon=False)
143 # hide tick and tick label of the big axes
144 plt.tick_params(labelcolor='none', top=False, bottom=False, left=
        False, right=False)
145 plt.grid(False)
146 x = np.linspace(0, 100, 100)
147 ax[0].plot(x, sample_params[:,0], "r+", label="Sampled $alpha$")
148 ax[1].plot(x, sample_params[:,1], "bo", label="Sampled $gamma_1$")
149 ax[2].plot(x, sample_params[:,2], "g*", label="Sampled $gamma_2$")
150 ax[3].plot(x, sample_params[:,3], "+", label="Sampled $f$")
151 plt.xlabel("Simulation (Individual Mouse)")
152 plt.ylabel("Sampled Parameter Value")
153 ax[0].legend(fontsize="8")
154 ax[1].legend(fontsize="8")
155 ax[2].legend(fontsize="8")
156 ax[2].legend(fontsize="8")
157 plt.show()
158 '''
159
160 # Calculate the correlation matrix - check for correlation between
        values
161 sampled_corr_matrix = np.corrcoef(sample_params, rowvar=False)
162 print(sampled_corr_matrix)
163
164 # Now create simulations of attacks and corresponding threat
        predictions for each set of params
165 sampled_attacks = np.zeros((100, 5400))

```



```

166 sampled_predictions = np.zeros((100, 5400))
167
168 for index, (a, g1, g2, f) in enumerate(zip(sample_params[:,0],
169     sample_params[:,1], sample_params[:,2], sample_params[:,3])):
170     # Generate day 1 sequence of attacks for each
171     A_shocks = day_1(15)
172     B_shocks = np.zeros_like(A_shocks)
173     # Stack for input to TD-Mom Model
174     contexts_stacked = np.column_stack((A_shocks, B_shocks))
175
176     # Generate threat predictions from this simulated sequence of
177     attacks
178     predictions = TD_momentum_model(contexts = contexts_stacked,
179     alpha = a, gamma1 = g1, gamma2 = g2, f = f)[0]
180
181     # Add to sampled attacks/preds arrays
182     sampled_attacks[index, :] = A_shocks
183     sampled_predictions[index, :] = predictions
184
185
186 np.sum(sampled_predictions, axis=1)
187 #np.where(np.sum(sampled_predictions, axis=1)==0)
188
189 # Fitting parameter values to these simulations (see if we can
190 recover params)
191 initial_vals = np.array((0.5, 0.95, 0.4, 1.5))
192
193 bounds = [(0.05, 1), (0.9, 1), (0, 0.8), (0, 3)]
194
195 opt_sampled_data = np.zeros((len(sampled_predictions[:,0]), 4))
196
197 # Minimization for simulations (Mice)
198 for m in range(0, len(sampled_predictions[:,0])):
199     # Compute minimization for each participant
200     B_shocks = np.zeros_like(sampled_attacks[m:m + 1].reshape(-1))
201     contexts_stacked = np.column_stack((sampled_attacks[m:m + 1].
202     reshape(-1), B_shocks))
203
204     result = differential_evolution(NLL_td_mom, x0=initial_vals,
205     bounds=bounds,
206
207                                     args=(contexts_stacked,
208     sampled_predictions[m:m + 1].reshape((5400, 1)), 0),

```

```
201         strategy='best1bin', polish=True
202     )
203     print(f'n_iter: {result.nit} - success: {result.success} - nll {
204         result.fun}')
205     alpha, gamma1, gamma2, f= result.x
206     # Add results to opt_params storing opt param values for each
207     participant
208     opt_sampled_data[m, 0] = alpha
209     opt_sampled_data[m, 1] = gamma1
210     opt_sampled_data[m, 2] = gamma2
211     opt_sampled_data[m, 3] = f
212 #Alpha params
213 multivariate_sim_alpha = sample_params[:,0]
214 fitted_simulated_alpha = opt_sampled_data[:,0]
215 alpha_pearson = scipy.stats.pearsonr(multivariate_sim_alpha,
216     fitted_simulated_alpha)
217 # gamma1 params
218 multivariate_sim_gamma1 = sample_params[:,1]
219 fitted_simulated_gamma1 = opt_sampled_data[:,1]
220 gamma1_pearson = scipy.stats.pearsonr(multivariate_sim_gamma1,
221     fitted_simulated_gamma1)
222 # gamma2 params
223 multivariate_sim_gamma2 = sample_params[:,2]
224 fitted_simulated_gamma2 = opt_sampled_data[:,2]
225 gamma2_pearson = scipy.stats.pearsonr(multivariate_sim_gamma2,
226     fitted_simulated_gamma2)
227 # f params
228 multivariate_sim_f = sample_params[:,3]
229 fitted_simulated_f = opt_sampled_data[:,3]
230 f_pearson = scipy.stats.pearsonr(multivariate_sim_f,
231     fitted_simulated_f)
```

D.5 TD-Momentum Model - List Sampled Parameter Recovery

```

1  ### DIFFERENTIAL EVOLUTION - LIST PARAMS ###
2  alpha_list = mat_contents["learning_rate_list"][0]
3  gamma1_list = mat_contents["decay_list"][0]
4  gamma2_list = mat_contents["momentum_rate_list"][0]
5  f_list = -mat_contents["scaling_list"][0]
6
7  sample_params = np.zeros((100, 4))
8
9  for s in range(100):
10     sample_params[s, 0] = np.random.choice(alpha_list)
11     sample_params[s, 1] = np.random.choice(gamma1_list)
12     sample_params[s, 2] = np.random.choice(gamma2_list)
13     sample_params[s, 3] = np.random.choice(f_list)
14
15  corr_matrix = np.corrcoef(sample_params, rowvar=False)
16  print(corr_matrix)
17
18  # Now create simulations of attacks and corresponding threat
19     predictions for each set of params
20
21  sampled_attacks = np.zeros((100, 5400))
22  sampled_predictions = np.zeros((100, 5400))
23
24  for index, (a, g1, g2, f) in enumerate(zip(sample_params[:,0],
25     sample_params[:,1], sample_params[:,2], sample_params[:,3])):
26     # Generate day 1 sequence of attacks for each
27     A_shocks = day_1(15)
28     B_shocks = np.zeros_like(A_shocks)
29     # Stack for input to TD-Mom Model
30     contexts_stacked = np.column_stack((A_shocks, B_shocks))
31
32     # Generate threat predictions from this simulated sequence of
33     attacks
34     predictions = TD_momentum_model(contexts = contexts_stacked,
35     alpha = a, gamma1 = g1, gamma2 = g2, f = f)[0]
36
37     # Add to sampled attacks/preds arrays
38     sampled_attacks[index, :] = A_shocks

```

```

35     sampled_predictions[index, :] = predictions
36
37 np.sum(sampled_predictions, axis=1)
38 # Low alpha and high f give large threat estimates
39
40 # Fitting parameter values to these simulations (see if we can
    recover params)
41
42 from scipy.optimize import differential_evolution
43 initial_vals = np.array((0.5, 0.95, 0.4, 1.5))
44
45 bounds = [(0.05, 1), (0.9, 1), (0, 0.8), (0, 3)]
46
47 opt_sampled_data = np.zeros((len(sampled_predictions[:,0]), 4))
48 # Minimization for simulations (Mice)
49 for m in range(0, len(sampled_predictions[:,0])):
50     # Compute minimization for each participant
51     B_shocks = np.zeros_like(sampled_attacks[m:m + 1].reshape(-1))
52     contexts_stacked = np.column_stack((sampled_attacks[m:m + 1].
    reshape(-1), B_shocks))
53
54     result = differential_evolution(NLL_td_mom, x0=initial_vals,
    bounds=bounds,
55                                     args=(contexts_stacked,
    sampled_predictions[m:m + 1].reshape((5400, 1)), 0),
56                                     strategy='best1bin', polish=True
    )
57
58     print(f'n_iter: {result.nit} - success: {result.success} - nll {
    result.fun}')
59     alpha, gamma1, gamma2, f= result.x
60     # Add results to opt_params storing opt param values for each
    participant
61     opt_sampled_data[m, 0] = alpha
62     opt_sampled_data[m, 1] = gamma1
63     opt_sampled_data[m, 2] = gamma2
64     opt_sampled_data[m, 3] = f
65
66 #Alpha params
67 selected_sim_alpha = sample_params[:,0]
68 fitted_simulated_alpha = opt_sampled_data[:,0]
69 alpha_pearson = scipy.stats.pearsonr(selected_sim_alpha,

```

```

    fitted_simulated_alpha)
70
71 # gamma1 params
72 selected_sim_gamma1 = sample_params[:,1]
73 fitted_simulated_gamma1 = opt_sampled_data[:,1]
74 gamma1_pearson = scipy.stats.pearsonr(selected_sim_gamma1,
    fitted_simulated_gamma1)
75
76 # gamma2 params
77 selected_sim_gamma2 = sample_params[:,2]
78 fitted_simulated_gamma2 = opt_sampled_data[:,2]
79 gamma2_pearson = scipy.stats.pearsonr(selected_sim_gamma2,
    fitted_simulated_gamma2)
80
81 # f params
82 selected_sim_f = sample_params[:,3]
83 fitted_simulated_f = opt_sampled_data[:,3]
84 f_pearson = scipy.stats.pearsonr(selected_sim_f, fitted_simulated_f)

```

D.6 BIC Model Comparison

```

1 # Unstressed model comparison: BIC_TD - BIC_TD-MOM
2
3 unstressed_comparison = pd.DataFrame({
4     "BIC_TD - BIC_TD-MOM": [None] * len(TD_Mom_unstressed_results)})
5
6 unstressed_comparison["BIC_TD - BIC_TD-MOM"] = TD_unstressed_results
    ["BIC Score"] - TD_Mom_unstressed_results["BIC Score"]
7
8
9 # Stressed model comparison: BIC_TD - BIC_TD-MOM
10 stressed_comparison = pd.DataFrame({
11     "BIC_TD - BIC_TD-MOM": [None] * len(TD_Mom_stressed_results)})
12
13 stressed_comparison["BIC_TD - BIC_TD-MOM"] = TD_stressed_results["
    BIC Score"] - TD_Mom_stressed_results["BIC Score"]
14
15 stressed_comparison["BIC_TD - BIC_TD-MOM"] = stressed_comparison["
    BIC_TD - BIC_TD-MOM"]*(-1)
16
17

```

```
18 ### Re-creating fig 5D: Model Comparison Plot
19 # Create a scatter plot
20 plt.figure(figsize=(8, 6))
21
22 # Plot unstressed_comparison on the top
23 plt.scatter(unstressed_comparison, range(len(unstressed_comparison))
24            , color='blue', label='Unstressed')
25
26 # Plot stressed_comparison on the bottom
27 plt.scatter(stressed_comparison, range(len(stressed_comparison)),
28            color='red', label='Stressed')
29
30 plt.axvline(x=0, color='black', linestyle='--')
31 plt.axhline(y=-1, color='black', linestyle='--')
32 plt.xlabel('$BIC_{TD} - BIC_{TD \ Momentum}$', fontsize=15)
33 plt.yticks([])
34 plt.title('Stressed vs. Unstressed Mice Model Comparison', fontsize
35          =15)
36 plt.legend()
37
38 # Add arrows with text annotations
39 plt.annotate('Favours TD Momentum', xy=(87, -1.7), xytext=(1, -2),
40            arrowprops=dict(facecolor='black', arrowstyle='simple')
41            , fontsize=15)
42
43 plt.annotate('Favours TD', xy=(-53, -1.7), xytext=(-35, -2),
44            arrowprops=dict(facecolor='black', arrowstyle='simple')
45            , fontsize=15)
46 plt.tight_layout()
47 plt.show()
```

Appendix E

Extension Simulations Code

E.1 Associability TD-Momentum Model

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3
4 def ASSOCIABILITY_TD_MOM(init, contexts, alpha, gamma1, gamma2, f,
5   eta):
6     '''
7     Evaluates TD Momentum model threat at each timestep in each
8     context only, no changing contexts, assumes the same
9     agent in one contex for the whole time, creates threat in other
10    context as momentum term only.
11
12    :param init: Provides initialisation points for threat readings
13    in both contexts (days 6 and 7)
14    :param contexts: Sequences of unconditioned stimuli across all
15    contexts. e.g. sefl(sims = 2)[1] for 2nd mouse sim
16    :param u: Key of chosen context to compute TD momentum model for
17    . (day1, day6, day7)
18    :param alpha: Learning rate
19    :param gamma1: Decay rate for threat
20    :param gamma2: Decay rate for momentum across all contexts
21    :param f: Scaling Constant for momentum term
22    :param eta: Associability weight parameter (between 0 and 1)
23
24    :return: TD Momentum model threat estimation levels over all
25    timesteps
26    '''
```

```

20     # Initialise momentum array
21     m = np.zeros_like(contexts[:,0])
22
23     T_A = np.zeros_like(contexts[:,0])
24     T_B = np.zeros_like(contexts[:,1])
25     Kappa_A = np.zeros_like(T_B)
26     Kappa_B = np.zeros_like(T_B)
27
28     # Set initialisation
29     T_A[0] = init[0]
30     T_B[0] = init[1]
31     Kappa_A[0] = 1
32     Kappa_B[0] = 1
33
34
35     for t in range(1, len(contexts[:,0])):
36         # Set PE to 0 at every time step before computing PE for
37         # current step
38         PE = 0
39         PE += alpha * (contexts[:, 0][t] - T_A[t - 1])
40         PE += alpha * (contexts[:, 1][t] - T_B[t - 1])
41
42         m[t] = m[t-1] + gamma2 * PE
43         m[t] = np.clip(m[t], 0, 1)
44
45         Kappa_A[t] = (1 - eta) * Kappa_A[t-1] + eta * (np.abs(
46             contexts[:, 0][t] - gamma1 * T_A[t - 1]))
47         Kappa_B[t] = (1 - eta) * Kappa_B[t-1] + eta * (np.abs(
48             contexts[:, 1][t] - gamma1 * T_B[t - 1]))
49
50         # Set lower bound constraint for associability values
51         if Kappa_A[t] < 0.05:
52             Kappa_A[t] = 0.05
53
54         if Kappa_B[t] < 0.05:
55             Kappa_B[t] = 0.05
56
57         T_A[t] = T_A[t - 1] + alpha * Kappa_A[t-1] * (contexts[:,
58             0][t] - gamma1 * T_A[t - 1]) + f * m[t]
59         T_B[t] = T_B[t - 1] + alpha * Kappa_B[t-1] * (contexts[:,
60             1][t] - gamma1 * T_B[t - 1]) + f * m[t]

```



```

57     T_A[t] = np.clip(T_A[t], 0, 1)
58     T_B[t] = np.clip(T_B[t], 0, 1)
59
60     #T_A, T_B = combined_rescale(T_A, T_B, 0, 1)
61     #Kappa_A, Kappa_B = combined_rescale(Kappa_A, Kappa_B, 0, 1)
62
63     return T_A, T_B, Kappa_A, Kappa_B
64
65
66
67 # Compare to this !
68 def TD_momentum_model(contexts, alpha, gamma1, gamma2, f):
69     '''
70     Evaluates TD Momentum model threat at each timestep in each
71     context only, no changing contexts, assumes the same
72     agent in one contex for the whole time, creates threat in other
73     context as momentum term only.
74
75     :param contexts: Sequences of unconditioned stimuli across all
76     contexts. e.g. sefl(sims = 2)[1] for 2nd mouse sim
77     :param u: Key of chosen context to compute TD momentum model for
78     . (day1, day6, day7)
79     :param alpha: Learning rate
80     :param gamma1: Decay rate for threat
81     :param gamma2: Decay rate for momentum across all contexts
82     :param f: Scaling Constant for momentum term
83
84     :return: TD Momentum model threat estimation levels over all
85     timesteps
86     '''
87     # Initialise momentum array
88     m = np.zeros_like(contexts[:,0])
89
90     T_A = np.zeros_like(contexts[:,0])
91     T_B = np.zeros_like(contexts[:,1])
92
93     for t in range(1, len(contexts[:,0])):
94         # Set PE to 0 at every time step before computing PE for
95         current step
96         PE = 0
97         PE += alpha * (contexts[:, 0][t] - T_A[t - 1])

```

```

93     PE += alpha * (contexts[:, 1][t] - T_B[t - 1])
94
95     m[t] = m[t-1] + gamma2 * PE
96     m[t] = np.clip(m[t], 0, 1)      # NEWLY ADDED 25/07/23
97
98     T_A[t] = T_A[t - 1] + alpha * (contexts[:, 0][t] - gamma1 *
99     T_A[t - 1]) + f * m[t]
100
101     T_B[t] = T_B[t - 1] + alpha * (contexts[:, 1][t] - gamma1 *
102     T_B[t - 1]) + f * m[t]
103
104     T_A[t] = np.clip(T_A[t], 0, 1)
105     T_B[t] = np.clip(T_B[t], 0, 1)
106
107     #T_A, T_B = combined_rescale(T_A, T_B, 0, 1)      # NEWLY ADDED
108     25/07/23
109
110     return T_A, T_B
111
112
113 # Create contexts attack sequences
114 #A = day_1(15)
115
116 A = np.zeros(1000)
117 # Insert 25 random 1s at random positions in the array
118 random_indices = np.random.choice(1000, 10, replace=False)
119 A[random_indices] = 1
120
121 B = np.zeros_like(A)
122 # Stack for input to TD-Mom Model
123 contexts_stacked = np.column_stack((A, B))
124 initial = [0, 0]
125
126 #x = np.linspace(0, 90, 162000)
127
128 x = np.linspace(0, 10, 1000)
129
130 y1, y2, y3, y4 = ASSOCIABILITY_TD_MOM(init = initial, contexts =
131     contexts_stacked, alpha =0.05, gamma1 = 0.9999,
132     gamma2 = 0.05, f = 0.1, eta = 0.01)

```

```
131 y5, y6 = TD_momentum_model(contexts = contexts_stacked, alpha =0.05,
    gamma1 = 0.9999,
132                               gamma2 = 0.05, f = 0.1)
133
134
135 fig, axes = plt.subplots(1, 3, figsize=(18, 6))
136
137 # Plot for the first graph
138 axes[1].plot(x, y1, label="A")
139 axes[1].plot(x, y2, label="B")
140 axes[1].plot(x, A, label="A input", alpha=0.20)
141 axes[1].plot(x, B, label="B input", alpha=0.20)
142 axes[1].set_title("Associability TD-Momentum Simulation ( $\eta =$ 
    0.01$)", fontsize=15)
143 axes[1].set_xlabel("Time", fontsize=15)
144 axes[1].set_ylabel("Threat", fontsize=15)
145 axes[1].set_ylim(0, np.max(y1))
146 axes[1].legend(loc="upper right")
147
148 # Plot for the second graph
149 axes[2].plot(x, y3, label="$\kappa_A$")
150 axes[2].plot(x, y4, label="$\kappa_B$")
151 axes[2].plot(x, A, label="A input", alpha=0.20)
152 axes[2].plot(x, B, label="B input", alpha=0.20)
153 axes[2].set_title("Associability Values", fontsize=15)
154 axes[2].set_xlabel("Time", fontsize=15)
155 axes[2].set_ylabel("Associability", fontsize=15)
156 axes[2].set_ylim(0, np.max(y3))
157 axes[2].legend(loc="upper right")
158
159 # Plot for the third graph
160 axes[0].plot(x, y5, label="A")
161 axes[0].plot(x, y6, label="B")
162 axes[0].plot(x, A, label="A input", alpha=0.20)
163 axes[0].plot(x, B, label="B input", alpha=0.20)
164 axes[0].set_xlabel("Time", fontsize=15)
165 axes[0].set_ylabel("Threat", fontsize=15)
166 axes[0].set_title("TD-Momentum Simulation", fontsize=15)
167 axes[0].set_ylim(0, np.max(y5))
168 axes[0].legend(loc="upper right")
169
170 plt.tight_layout()
```

```
171 plt.show()
```

E.2 Risk-Sensitive TD-Momentum Model

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3
4 def RISK_SENSITIVE_TD_MOM(init, contexts, alpha_pos, alpha_neg,
5   gamma1, gamma2, f):
6     '''
7     Evaluates TD Momentum model threat at each timestep in each
8     context only, no changing contexts, assumes the same
9     agent in one context for the whole time, creates threat in other
10    context as momentum term only.
11
12    :param init: Provides initialisation points for threat readings
13    in both contexts (days 6 and 7)
14    :param contexts: Sequences of unconditioned stimuli across all
15    contexts. e.g. sefl(sims = 2)[1] for 2nd mouse sim
16    :param alpha_pos: Learning rate for positive prediction error
17    :param alpha_neg: Learning rate for negative prediction error
18    :param gamma1: Decay rate for threat
19    :param gamma2: Decay rate for momentum across all contexts
20    :param f: Scaling Constant for momentum term:
21
22    :return: TD Momentum model threat estimation levels over all
23    timesteps
24    '''
25    # Initialise momentum array
26    m = np.zeros_like(contexts[:,0])
27
28    T_A = np.zeros_like(contexts[:,0])
29    T_B = np.zeros_like(contexts[:,1])
30
31    # Set initialisation
32    T_A[0] = init[0]
33    T_B[0] = init[1]
34
35    for t in range(1, len(contexts[:,0])):
36        # Selecting learning rate for time step:
37        # For context A:

```

```

32     if (contexts[:, 0][t] - gamma1 * T_A[t - 1]) < 0:
33         alpha_A = alpha_neg
34
35     if (contexts[:, 0][t] - gamma1 * T_A[t - 1]) >= 0:
36         alpha_A = alpha_pos
37
38     # For context B:
39     if (contexts[:, 1][t] - gamma1 * T_B[t - 1]) < 0:
40         alpha_B = alpha_neg
41
42     if (contexts[:, 1][t] - gamma1 * T_B[t - 1]) >= 0:
43         alpha_B = alpha_pos
44
45
46     # Set PE to 0 at every time step before computing PE for
47     current step
48     PE = 0
49     PE += alpha_A * (contexts[:, 0][t] - gamma1 * T_A[t - 1])
50     PE += alpha_B * (contexts[:, 1][t] - gamma1 * T_B[t - 1])
51
52     m[t] = m[t-1] + gamma2 * PE
53     m[t] = np.clip(m[t], 0, 1)
54
55     T_A[t] = T_A[t - 1] + alpha_A * (contexts[:, 0][t] - gamma1
56 * T_A[t - 1]) + f * m[t]
57     T_B[t] = T_B[t - 1] + alpha_B * (contexts[:, 1][t] - gamma1
58 * T_B[t - 1]) + f * m[t]
59
60     T_A[t] = np.clip(T_A[t], 0, 1)
61     T_B[t] = np.clip(T_B[t], 0, 1)
62
63     #T_A, T_B = combined_rescale(T_A, T_B, 0, 1)
64
65     return T_A, T_B
66
67 # Compare to this !
68 def TD_momentum_model(contexts, alpha, gamma1, gamma2, f):
69     '''
70     Evaluates TD Momentum model threat at each timestep in each

```

```

context only, no changing contexts, assumes the same
71 agent in one contex for the whole time, creates threat in other
context as momentum term only.
72
73 :param contexts: Sequences of unconditioned stimuli across all
contexts. e.g. sefl(sims = 2)[1] for 2nd mouse sim
74 :param u: Key of chosen context to compute TD momentum model for
. (day1, day6, day7)
75 :param alpha: Learning rate
76 :param gamma1: Decay rate for threat
77 :param gamma2: Decay rate for momentum across all contexts
78 :param f: Scaling Constant for momentum term
79
80 :return: TD Momentum model threat estimation levels over all
timesteps
81 '''
82 # Initialise momentum array
83 m = np.zeros_like(contexts[:,0])
84
85 T_A = np.zeros_like(contexts[:,0])
86 T_B = np.zeros_like(contexts[:,1])
87
88
89 for t in range(1, len(contexts[:,0])):
90     # Set PE to 0 at every time step before computing PE for
current step
91     PE = 0
92     PE += alpha * (contexts[:, 0][t] - T_A[t - 1])
93     PE += alpha * (contexts[:, 1][t] - T_B[t - 1])
94
95     m[t] = m[t-1] + gamma2 * PE
96     m[t] = np.clip(m[t], 0, 1) # NEWLY ADDED 25/07/23
97
98     T_A[t] = T_A[t - 1] + alpha * (contexts[:, 0][t] - gamma1 *
T_A[t - 1]) + f * m[t]
99     T_B[t] = T_B[t - 1] + alpha * (contexts[:, 1][t] - gamma1 *
T_B[t - 1]) + f * m[t]
100
101     T_A[t] = np.clip(T_A[t], 0, 1)
102     T_B[t] = np.clip(T_B[t], 0, 1)
103
104 #T_A, T_B = combined_rescale(T_A, T_B, 0, 1) # NEWLY ADDED

```

```
25/07/23
105
106     return T_A, T_B
107
108
109
110 # Create context attack sequences
111 A = np.zeros(1000)
112 # Insert 25 random 1s at random positions in the array
113 random_indices = np.random.choice(1000, 10, replace=False)
114 A[random_indices] = 1
115
116 B = np.zeros_like(A)
117 # Stack for input to TD-Mom Model
118 contexts_stacked = np.column_stack((A, B))
119 initial = [0, 0]
120
121 #x = np.linspace(0, 90, 162000)
122
123 x = np.linspace(0, 10, 1000)
124
125 y1, y2 = RISK_SENSITIVE_TD_MOM(init = initial, contexts =
    contexts_stacked, alpha_pos = 0.05, alpha_neg = 0.2, gammal =
    0.9999,
126                                     gamma2 = 0.05, f = 0.1)
127
128
129 y3, y4 = TD_momentum_model(contexts = contexts_stacked, alpha =0.05,
    gammal = 0.9999,
130                                     gamma2 = 0.05, f = 0.1)
131
132
133 fig, axes = plt.subplots(1, 2, figsize=(12, 6))
134
135 # Plot for the first graph
136 axes[0].plot(x, y3, label="A")
137 axes[0].plot(x, y4, label="B")
138 axes[0].plot(x, A, label="A input", alpha=0.20)
139 axes[0].plot(x, B, label="B input", alpha=0.20)
140 axes[0].set_xlabel("Time", fontsize=15)
141 axes[0].set_ylabel("Threat", fontsize=15)
142 axes[0].set_title("TD-Momentum Simulation", fontsize=15)
```

```

143 axes[0].set_ylim(0, np.max(y3))
144 axes[0].legend(loc="upper right")
145
146 # Plot for the second graph
147 axes[1].plot(x, y1, label="A")
148 axes[1].plot(x, y2, label="B")
149 axes[1].plot(x, A, label="A input", alpha=0.20)
150 axes[1].plot(x, B, label="B input", alpha=0.20)
151 axes[1].set_title("Risk-Sensitive TD-Momentum Simulation", fontsize
    =15)
152 axes[1].set_xlabel("Time", fontsize=15)
153 axes[1].set_ylabel("Threat", fontsize=15)
154 axes[1].set_ylim(0, np.max(y1))
155 axes[1].legend(loc="upper right")
156
157 plt.tight_layout()
158 plt.show()

```

E.3 Valence-Partitioned TD-Momentum Model

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3
4
5 def VP_TD_MOM(init, contexts, alpha_pos, alpha_neg, gamma1, gamma2,
6 f):
7     '''
8     Evaluates TD Momentum model threat at each timestep in each
9     context only, no changing contexts, assumes the same
10    agent in one context for the whole time, creates threat in other
11    context as momentum term only.
12
13    :param init: Provides initialisation points for threat readings
14    in both contexts (days 6 and 7)
15    :param contexts: Sequences of unconditioned stimuli across all
16    contexts. e.g. sefl(sims = 2)[1] for 2nd mouse sim
17    :param alpha_pos: Learning rate for positively valenced outcomes
18    :param alpha_neg: Learning rate for negatively valenced outcomes
19    :param gamma1: Decay rate for threat
20    :param gamma2: Decay rate for momentum across all contexts
21    :param f: Scaling Constant for momentum term:

```



```

17
18     :return: TD Momentum model threat estimation levels over all
19     timesteps
20     '''
21     # Initialise momentum arrays
22     m_P = np.zeros_like(contexts[:,0])
23     m_N = np.zeros_like(m_P)
24
25     # Overall Threat in each context
26     T_A = np.zeros_like(contexts[:,0])
27     T_B = np.zeros_like(contexts[:,1])
28
29     # Partitioned values for each context
30     VA_P = np.zeros_like(T_A)
31     VA_N = np.zeros_like(T_A)
32
33     VB_P = np.zeros_like(T_B)
34     VB_N = np.zeros_like(T_B)
35
36     # Set initialisation
37     T_A[0] = init[0]
38     T_B[0] = init[1]
39
40     for t in range(1, len(contexts[:,0])):
41         # Partitioning based on outcome for each context:
42         # For context A:
43         if contexts[:, 0][t] > 0:
44             A_delta_P = contexts[:, 0][t] - gamma1 * VA_P[t - 1]
45
46         if contexts[:, 0][t] <= 0:
47             A_delta_P = 0 - gamma1 * VA_P[t - 1]
48
49         if contexts[:, 0][t] < 0:
50             A_delta_N = np.abs(contexts[:, 0][t]) - gamma1 * VA_N[t
51 - 1]
52
53         if contexts[:, 0][t] >= 0:
54             A_delta_N = 0 - gamma1 * VA_N[t - 1]
55
56         # For context B:
57         if contexts[:, 1][t] > 0:

```

```

57         B_delta_P = contexts[:, 1][t] - gamma1 * VB_P[t - 1]
58
59         if contexts[:, 1][t] <= 0:
60             B_delta_P = 0 - gamma1 * VB_P[t - 1]
61
62         if contexts[:, 1][t] < 0:
63             B_delta_N = np.abs(contexts[:, 1][t]) - gamma1 * VB_N[t
64 - 1]
65
66         if contexts[:, 1][t] >= 0:
67             B_delta_N = 0 - gamma1 * VB_N[t - 1]
68
69
70         # Set PE to 0 at every time step before computing PE for
71         # current step
72         PE_P = 0
73         PE_P += alpha_pos * A_delta_P
74         PE_P += alpha_pos * B_delta_P
75
76         m_P[t] = m_P[t-1] + gamma2 * PE_P
77         #m_P[t] = np.clip(m_P[t], 0, 1) DO WE NEED THESE? BETWEEN -1
78         AND 1?
79
80         PE_N = 0
81         PE_N += alpha_neg * A_delta_N
82         PE_N += alpha_neg * B_delta_N
83
84         m_N[t] = m_N[t-1] + gamma2 * PE_N
85         #m_P[t] = np.clip(m_P[t], 0, 1) DO WE NEED THESE? BETWEEN -1
86         AND 1?
87
88         # Positive and negative value functions for each context
89         VA_P[t] = VA_P[t - 1] + alpha_pos * A_delta_P + f * m_P[t]
90         VA_N[t] = VA_N[t - 1] + alpha_neg * A_delta_N + f * m_N[t]
91
92         VB_P[t] = VB_P[t - 1] + alpha_pos * B_delta_P + f * m_P[t]
93         VB_N[t] = VB_N[t - 1] + alpha_neg * B_delta_N + f * m_N[t]
94
95         # Generating overall threat prediction for each context

```



```

133
134
135
136 # VP Plot
137 fig, axes = plt.subplots(1, 1, figsize=(6, 3))
138
139 # Plot for the first graph
140 axes.plot(x, y1, label="A")
141 axes.plot(x, y2, label="B")
142 axes.plot(x, A, label="A input", alpha=0.20)
143 axes.plot(x, B, label="B input", alpha=0.20)
144 axes.set_title("Valence-Partitioned TD-Momentum Simulation",
145               fontsize=15)
146 axes.set_xlabel("Time", fontsize=15)
147 axes.set_ylabel("Threat", fontsize=15)
148 axes.set_ylim(np.min(y1), np.max(y1))
149 axes.legend(loc="upper right")
150
151 plt.tight_layout()
152 plt.show()
153
154 # momentum plot
155 fig, axes = plt.subplots(1, 1, figsize=(6, 3))
156
157 # Plot for the first graph
158 axes.plot(x, mp, label="$m^P_t$")
159 axes.plot(x, mn, label="$m^N_t$")
160 axes.plot(x, A, label="A input", alpha=0.20)
161 axes.plot(x, B, label="B input", alpha=0.20)
162 axes.set_title("Valence-Partitioned Momentum Terms", fontsize=15)
163 axes.set_xlabel("Time", fontsize=15)
164 axes.set_ylabel("Momentum Value", fontsize=15)
165 axes.set_ylim(np.min(y1), np.max(y1))
166 axes.legend(loc="upper right")
167
168 plt.tight_layout()
169 plt.show()
170
171 # Positive and negative valence functions for each context
172 fig, axes = plt.subplots(1, 2, figsize=(12, 6))
173

```

```
174 # Plot for the first graph
175 axes[0].plot(x, y5, label="$V^P_A$")
176 axes[0].plot(x, y6, label="$V^N_A$")
177 axes[0].plot(x, A, label="A input", alpha=0.20)
178 axes[0].set_xlabel("Time", fontsize=15)
179 axes[0].set_ylabel("Threat", fontsize=15)
180 axes[0].set_title("Valence-Partitioned Value Functions - Context A",
181                  fontsize=15)
182 #axes[0].set_ylim(0, np.max(y3))
183 axes[0].legend(loc="upper right")
184
185 # Plot for the second graph
186 axes[1].plot(x, y7, label="$V^P_B$")
187 axes[1].plot(x, y8, label="$V^N_B$")
188 axes[1].plot(x, B, label="B input", alpha=0.20)
189 axes[1].set_xlabel("Time", fontsize=15)
190 axes[1].set_ylabel("Threat", fontsize=15)
191 axes[1].set_title("Valence-Partitioned Value Functions - Context B",
192                  fontsize=15)
193 #axes[1].set_ylim(np.min(y1), np.max(y1))
194 axes[1].legend(loc="upper right")
195
196 plt.tight_layout()
197 plt.show()
```