



Kostova, Ralitsa Angelova (2023) *Neural and computational principles of social vs non-social decision-making under risk*. PhD thesis.

<http://theses.gla.ac.uk/83943/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

NEURAL AND
COMPUTATIONAL
PRINCIPLES OF SOCIAL VS
NON-SOCIAL
DECISION-MAKING UNDER
RISK

Ralitsa Angelova Kostova, BSc, MSc

SUBMITTED IN FULFILMENT OF THE REQUIREMENTS FOR THE
DEGREE OF
DOCTOR OF PHILOSOPHY

SCHOOL OF PSYCHOLOGY AND NEUROSCIENCE
COLLEGE OF SCIENCE AND ENGINEERING



University
of Glasgow

06 2023

Abstract

Decisions feel daunting, as the weight of selecting one path over another engenders a sense of unease and hesitation. This is a natural consequence of the fact that most decisions in the real world involve uncertainty. A choice involves two or more alternatives and usually resolves in the selection of a subjectively preferable option. Very often, however, decision-making requires consideration of multiple options and their numerous possible outcomes, as well as determining what 'preferable' stands for. This makes choices 'risky.' With this thesis, I studied the neural principles associated with risk, and I tested fluctuations of risk-taking as predicted by a novel computational model and social identity theory. In the first experiment, I used social stimuli as cues and recorded trial-by-trial fluctuations in EEG to try and capture brain responses to estimates of risk and risk prediction errors. The results reveal distinct spatio-temporal EEG components associated with the computation of risk and violations of expected risk. In the second experiment, I tested mediators of trial-by-trial risk-seeking. More specifically, independent of participants' risk propensities in real life, I implemented a task to drive risk-taking choice in some trials and risk avoidance in others. I showed that in non-social contexts positive reward prediction errors can predict risk-taking, and I found brain responses associated with this process. In the final chapter, I discuss the same task in which I tried to induce risk-seeking with the addition of a social factor aiming to test predictions from social identity theory. I showed that the mere online presence of an in-group and an out-group member was enough to alter behaviour during the task, although possible explanations can span from exploratory behaviour in some groups of participants, to overall arousal or increased stress in the participants while being observed. Together these experiments show the importance of incorporat-

ing social factors into studies of decision-making, the benefit of computational methods for a better understanding of risky decision-making and model-based neural responses, and the importance of accounting for individual differences when studying value-based choice.

Contents

Abstract	ii
Acknowledgements	xii
Declaration	xiv
Abbreviations	xv
1 Background	1
1.1 General decision-making	3
1.2 Learning and value	5
1.2.1 Reinforcement learning	5
1.2.2 Reinforcement learning and dopamine	8
1.2.3 Basal ganglia, rewards, and learning	9
1.3 Risky and social decision-making	12
1.3.1 Defining risk	12
1.3.2 Human studies with mean-variance dissociation	17
1.3.3 The basal ganglia model of risk-taking	20
1.3.4 Social decision-making, risk, and learning	23
1.3.5 Social group membership, cortico-striatal inputs, and risk-seeking	29
1.4 The goals of this thesis	31
2 Spatiotemporal Representation of Risk and Risk Prediction Errors with Social Stimuli	34
2.1 Background	35
2.2 Materials and Methods	38
2.2.1 Participants	38

2.2.2	Risk and RiPE estimation	39
2.2.3	Stimuli and task	43
2.2.4	Main task	45
2.2.5	EEG data acquisition	49
2.2.6	EEG data pre-processing	49
2.2.7	Single-trial EEG analysis	50
2.2.8	Bootstrapping-based exclusion	52
2.2.9	Identifying risk-related EEG components	52
2.3	Results	56
2.3.1	Current risk	57
2.3.2	Absolute RiPE	57
2.3.3	Signed RiPE	58
2.4	Discussion	61

3 Risk-Seeking Driven by Stimulus Prediction Errors: Modelling and EEG Results **66**

3.1	Background	67
3.2	Methods	74
3.2.1	Participants	74
3.2.2	Task	75
3.2.3	Theoretical analysis	78
3.2.4	Additional materials	84
3.2.5	Procedure	84
3.2.6	Data analysis	85
3.3	Results	87
3.3.1	Behavioural results	88
3.3.2	Model comparisons	92
3.3.3	PEIRS2-stimulus-PEs predicting risk-taking	96
3.3.4	Model-based EEG results. Signed Stimulus-PE	97
3.3.5	Predicting risk-taking from EEG amplitudes at signed stimulus-PE100	
3.3.6	Exploratory analysis	101

3.3.7	Outcome-PEs	102
3.3.8	Relating stimulus-PEs to outcome-PEs	103
3.4	Discussion	104
3.4.1	Cause of risk-taking	104
3.4.2	Spatio-temporal brain patterns of stimulus-PEs	108
3.4.3	Prediction errors post-stimulus and post-outcome	110
3.4.4	Limitatons and future directions	112
4	Group Membership Bias and Stimulus Prediction Errors Influence on Risk-Seeking	114
4.1	Background	115
4.2	Methods	119
4.2.1	Social conditions	119
4.2.2	EEG analysis	123
4.2.3	Behavioural analysis	123
4.3	Results	124
4.3.1	Behavioural results	125
4.3.2	Model selection	130
4.3.3	Model-based EEG results	131
4.3.4	Exploratory Anxiety effects	134
4.3.5	Response to wins and losses. Outcome-PEs in relation to social condition	138
4.4	Discussion	140
4.4.1	Social bias or exploration	140
4.4.2	Risk-taking under stress, arousal, and anxiety	143
4.4.3	Limitations and future directions	146
5	General Discussion	148
5.1	Risk-processing and risk-seeking in social contexts	148
5.2	Encoding predicted risk separately from risk prediction errors	149
5.3	Risk-seeking influenced by positive reward prediction errors	152

5.4	Presence of others and anxiety on risky decision-making	154
5.4.1	Implications for addiction and gambling	156
5.5	Limitations and future directions	157
5.6	Conclusion	158

Bibliography		160
---------------------	--	------------

List of Tables

2.1	Computed values of each risk condition across trials. Shading represents LDA training and testing levels. Blue shading shows each condition's lowest (light blue) and highest (darker blue) values.	42
-----	---	----

List of Figures

1.1	A representation of different stages of perceptual and value-based decision-making	5
1.2	Dopaminergic structures associated with decision-making.	10
1.3	Risky and social decision-making.	12
1.4	Risk and reward relationship.	18
1.5	A schematic representation of long-term potentiation and depression. . . .	22
2.1	Values of pairs of stimuli.	41
2.2	Correlation plot of computed values.	41
2.3	Computed levels of risky and safe trials.	43
2.4	Example stimuli.	45
2.5	Experimental timeline.	48
2.6	Example of results of k-means clustering of average scalp topographies. . .	54
2.7	EEG-derived values of each risk variable across the four levels of trials. . .	56
2.8	Average discrimination ($n = 40$) and average topographies from best discrimination points.	59
2.9	EEG amplitudes at all three conditions derived from trials used for training the LDA (pink) and testing the resulting weights (black).	59
2.10	Density plots of the time of onset of best points relative to the onset of the stimulus.	60
3.1	Theoretical influence of situational stimulus PEs on cortico-striatal activity. .	69
3.2	Dopamine influence on risk-taking.	71
3.3	Performance on the behavioural task across participants.	74
3.4	Four example slot machines with their associated reward distributions. . .	76
3.5	Timeline of the bandit task.	77

3.6	A schematic of two instances of slot-machine combinations and the proposed basal ganglia mechanism for each one (modified from Moeller et al., 2021).	79
3.7	Risky choices for both-high and both-low options.	88
3.8	Learning curve for risk-taking during both-high and both-low combinations.	89
3.9	Reaction times to choose an option when presented with each pair of stimuli.	90
3.10	Pairs plot between risk-taking for each condition, anxiety scales, and risk-propensity scale.	91
3.11	Model comparisons for the baseline condition. BIC and model-simulated data.	94
3.12	Model comparisons for the baseline condition - 2.	95
3.13	Discriminator performance.	98
3.14	Scalp topography from best peaks averaged across participants.	99
3.15	Scalp topographies from best peaks from a sample of participants.	99
3.16	Y values (in micro-volts) plotted for each set of stimulus-PEs.	100
3.17	Binned Y values in ascending order with boxplots of the probability of risky choices.	101
3.18	Az (leave-one-out) discriminator performance for Signed and Absolute outcome-PEs and topographic maps averaged across peak times.	102
3.19	Correlations of topographic maps.	103
4.1	Image shown to the participants to assign them into social groups.	120
4.2	An example screen of a pre-recorded Zoom connection with a confederate.	122
4.3	Ratings on belief that the in-group and out-group members observed the experiment.	125
4.4	Probability of risk-taking in both-high and both-low trials across the social blocks.	126
4.5	Learning curves for the social conditions.	127
4.6	Reaction times are shown as average values per pair of expected values for each social condition.	128
4.7	Average performance across blocks with standard error of the mean.	129
4.8	Model comparisons for the full task.	131

4.9	Discriminator performance and best maps.	132
4.10	Average Y values for trained and tested (grey) trials.	133
4.11	Correlation plot of risk-taking during baseline and social conditions and trait (TAI) and state (SAI) anxiety.	135
4.12	Percentage correct choices (of high EV within the different-means condition) across blocks.	137
4.13	Average Y activation across blocks, anxiety levels, and EV-pairs for signed stimulus-PEs.	138
4.14	Boxplots of maximum Az values between 150 ms and 350 ms post-feedback onset for the two social blocks.	139

Acknowledgements

First and foremost, I want to thank my supervisor Prof Marios G. Philiastides for the opportunity to embark on the PhD quest that had the misfortune to start with a world pandemic. Thank you, Marios, for always being calm and patient with all the challenges imposed by the circumstances and my learning process. Thank you for striking the best balance between support and autonomy, and for focusing on collaboration rather than imposing a hierarchy, helping me grow and learn by my own means. Thank you for uplifting my confidence when I felt doubt and for all the fruitful, intellectual, and stimulating conversations we had. Your professionalism, integrity, and dedication to academic excellence have resonated with what I always valued most in research, and why I commenced this PhD in the first place.

Thank you to all my previous supervisors who took part in my development as a researcher. Thank you Simon Garrod, Lawrence Barsalou, and Peter Uhlhaas. Special thank you to Dr Sara Sereno and Prof Edwin Robertson for the invaluable support and suggestions throughout the years of my PhD.

Thank you, Candice C. Morey for your brief but extremely valuable supervisions. You are truly an inspirational role model.

I want to thank my parents for helping me go to a great university and follow my passion.

Татяна, Ангел, и Илиана, благодаря ви за безспирната подкрепа и оптимизъм. Вие сте най-доброто семейство, което човек може да иска.

I would like to thank the people who have been by my side all 10 years of my education. Thank you, my wonderful girls, Zlatina, Nikoleta, Dimana. Thank you, Vladimir, Theodor, Niya, Simeon, Petar, Martin, Desi, for carrying a bit of home and sharing it. Thank you Alex V. for inspiring me to study abroad. And special thanks to Nikolay for the never-ending philosophical discourse and wonder about the human condition. Special thanks to Maryana Dancheva and all other high school teachers who had an extremely positive effect on my personal development.

I also want to thank the people who were mostly part of the end of my PhD and made everything full of life and enjoyment: Bianca, Laura, Eleonora, Jelena, Joana, Sean, Sander, Mircea, Luca, Christopher, Christoph. Thank you, Madeline, for being a mentee and a mentor for the last part of it. And thank you, Alex M.. for sharing this love for dopamine, empathy, and for helping me have a calm and smooth ride.

Declaration

I declare that, except where explicit reference is made to the contribution of others, that this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Ralitsa Angelova Kostova

Abbreviations

- ACC - Anterior Cingulate Cortex
- AUC - Area Under the Curve
- CS - Conditioned Stimulus
- DA - Dopamine
- dmPFC - dorso-medial Prefrontal Cortex
- EEG - Electroencephalogram
- EV - Expected Value
- fMRI - functional Magnetic Resonance Imaging
- IFG - Inferior Frontal Gyrus
- LDA - Linear Discriminant Analysis
- OFC - Orbito-frontal cortex
- PE - Prediction Error
- PCA - Principal Component Analysis
- RL - Reinforcement Learning
- RPE - Reward Prediction Error
- RiPE - Risk Prediction Error
- TD - Temporal Difference (learning)
- US - Unconditioned Stimulus
- VTA - Ventral Tagmental Area
- vmPFC - ventro-medial Prefrontal Cortex
- SD - Standard Deviation
- SN - Substantia Nigra

Chapter 1

Background

To understand risky decision-making in social and non-social settings, we need to understand the process of decision-making in general, what it encapsulates, how it can be studied and then consider whether the same process generalises across contexts. In its basic form in nature, decision-making entails some kind of action towards or away from an object. Whether this is to obtain food or escape from a predator, better decisions mean better chances of survival. In both animals and humans, better decisions come with experience and the ability to deliberate about costs and benefits. As such, decision-making processes are fundamentally intertwined with learning. It would be futile if every choice is deliberated on the spot without any prior information or memory. Subsequently, often instead of contemplating every decision, organisms use innate impulses, generalise across situations and stimuli, form biases, and use heuristics to derive the best course of action. Additionally, one can expect that simpler organisms, like insects for example, do not possess cognitive abilities that let them deliberate their choices, hence, would follow a more simplistic model of behaviour allowing the study of fundamental processes and predictions. Most of these processes can be modelled mathematically in humans and are theorised to form the basis for a completely rational decision-maker, the ‘homo-economicus.’ However, in the following pages, we will see that decision-making by humans involves both rationalisations of choices, as predicted by economics, but also emotional and social factors.

The realisation by economists that human choices: 1) cannot be predicted based on objective variables, 2) are not the same given the same circumstances, and 3) are bound to emotions, formed the basis of behavioural economics (Kahneman and Tversky, 1979). On the other hand, with the advances in brain imaging, computational methods, and findings from animal studies, neuroscience has contributed separately to the understanding of decision-making. Most recently, scientists from the fields of psychology, economics, and neuroscience have focused on merging the separate theories and findings of each field into a new paradigm of studying human behaviour, namely, neuroeconomics. The current work rests on the shoulder of an era of research trying to devise the best theories to explain value, independent of whether it comes from primary rewards in animals such as food, or secondary rewards in humans such as social approval. This multidisciplinary effort has produced many compelling computational models and testable hypotheses which informed our experimental design, modelling of behaviour and interpretation of electrophysiological results. I show results that support emerging theories about short-term risk-seeking influenced by positive prediction errors. This is consistent with theories about the interaction of networks in the basal ganglia and prefrontal cortex. Further, the results revealed the spatio-temporal pattern in the brain that encodes the different prediction errors leading to a choice. At the point of incorporating a social component into the experiment, I stumbled upon interesting findings that can be explained by an interplay of social biases and dispositional traits. The complexity of the behavioural and electrophysiological results points to the important task of future decision-making experiments to incorporate measures of individual differences and variable social conditions, as well as aiming for an interdisciplinary approach.

I based one of the experiments on a theory that proposes a division of risk into separate components that can be studied independently from reward. The three risk components that I found are suggested to drive learning under uncertainty and to be generalised in the brain for social and non-social stimuli. A challenge posed by the design of the experiment was the imminent trade-off between control over a study and its ecological validity. Although I highly controlled the relationship between variables and the allowed

behaviour in order to reduce the possibility of confounding factors and unexpected brain responses, a lot of the natural vigour and motivation in the participants to attend to a decision-making situation was sacrificed. Future studies aiming to reduce the decision-making process to its base components should consider the importance of the complex interplay between decision-making and motivation, attention, motor activation, cost of effort, and related elements of behaviour.

1.1 General decision-making

Decision-making involves a few necessary stages to be complete, and to be informative and functional for the organism. It also involves many brain areas that interact to interpret sensory inputs, incorporate memories of previous experiences, and prepare a motor response that is aligned with the best course of action. These stages have been conceptualised as the four-module decision-making process (Gold and Shadlen, 2007; Heekeren et al., 2008; Rangel et al., 2008), with *module* being a separate unit with a specialised function and occurring at a different stage during each decision-making episode. These modules, also shown on Figure 1.1 are evidence representation, decision-making, action selection, and outcome evaluation. When presented with a choice between two or more options, an organism forms a representation of the alternatives, which could be between objects moving left versus right (perceptual) or, for example, food such as chocolate versus broccoli (value-based). The second stage encompasses decision-making, that is the module determining the level of uncertainty of the evidence representation, dictating if the evidence is enough and if more attention is required. This decision is weighted by information gathered from previous experience if any exists, and current representations. After the first two stages are complete, it is time for action selection, most strongly related to an initiated motor response that already integrates a decision variable formed during the first two stages. The decision variable dictates the preference for one choice over another and drives motor response towards the relevant stimulus. Although decision-making may seem complete by now, the first three stages can be

utilised to extract information, adapt, and prepare for future similar decisions. This leads to the fourth stage, outcome evaluation, which involves the detection of errors conditional on provided external feedback or the fulfilment of an internal criterion. This stage directly influences learning and the adjustment of the other modules. Although conceptualising decision-making as a process with clear-cut stages facilitates testing, it is accepted that stages can be shuffled. For example, valuation may come before action selection or in parallel (Rangel et al., 2008). The four-stage framework was initially developed for perceptual decision-making and later adopted by theories of value-based decision-making with the appropriate re-conceptualisation of every stage relevant to values as shown on Figure 1.1.

Value-based decision-making deals with choices between alternatives towards which an organism can have different subjective interpretations and preferences. While perceptual decision-making has been largely informed by psychophysics and psychometric curves, and it can be predicted mainly from the properties of external stimuli, the science of value-based decision-making does not have the privilege of objective measures of preferences apart from the observed choice made by an organism. Although we can expect that reducing the contrast of a moving object would reduce the ability of any tested animal to distinguish the direction of movement (perceptual), we cannot say for certain that if one monkey prefers a banana to a strawberry, another monkey would show the same preference. This is why it is important to make clear what the concepts of value, reward, and punishment entail, and how they can be operationalised to be tested and measured within brain activity.

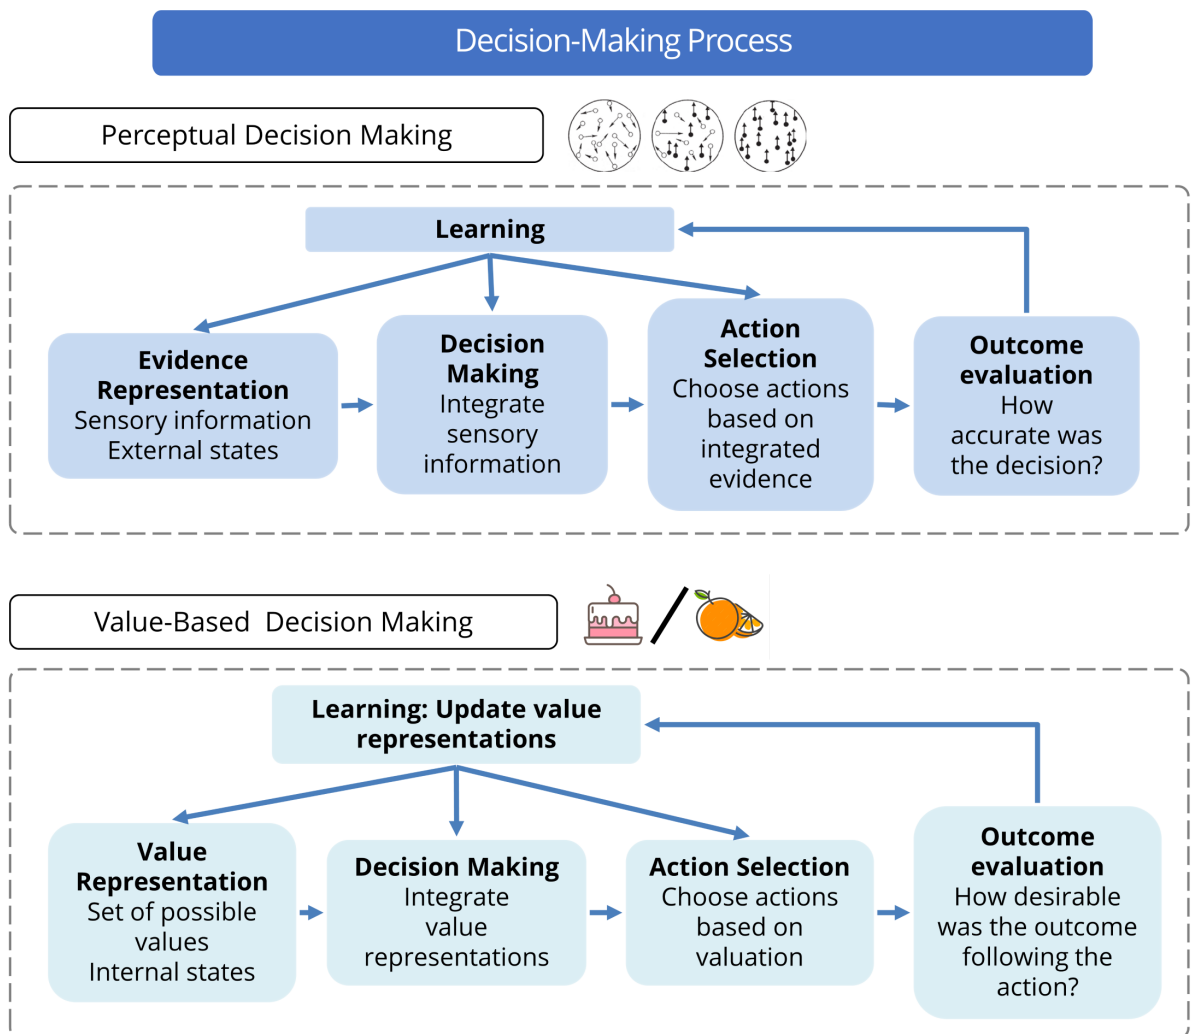


Figure 1.1 A representation of different stages of perceptual and value-based decision-making (modified from Rangel et al., 2008).

1.2 Learning and value

1.2.1 Reinforcement learning

Value is a determinant of how much something is worth, it is a concept encapsulating the subjective system by which some object in the environment or action is regarded as beneficial or costly for the organism. We expect organisms to prioritise options that have a high subjective value. However, unlike perceptual decision-making, studying

value-based decision-making poses the additional necessity to learn and then infer what is a high-valued or a low-valued option for the studied human or non-human participant. Moreover, there is an ensuing complexity in the valuation process that is specific to the organism, the situation, and the conditions during which decisions are made.

With the aim of understanding value-based decision-making, studies using animals have started working with the simplest values, which are defined as primary rewards, or primary reinforcers. At the beginning of the 1900s, Ivan Pavlov used dogs as test subjects and paired the presentation of a primary reward (food) with the ring of a bell (Akpan, 2020). The automatic response to food was increased salivation, while the ring of the bell was defined as a neutral stimulus as it did not produce any particular response in dogs. He discovered that after a few repetitions of presenting a neutral stimulus, the bell, right before the delivery of food led to salivation even when it was rung without the presence of food. The response to the unconditioned stimulus (food) was transferred as well to the now conditioned stimulus (the bell), forming an ‘associate strength.’ Shifting the focus from the passive reception of rewards to active behaviour in response to a stimulus, Thorndike’s work in the 1800s laid the foundation for the development of the Law of Effect, which states that behaviours that lead to positive outcomes are more likely to be repeated. Later, B.F. Skinner showed that animals learn to associate certain behaviours with either reward or punishment, which in turn can lead to either reinforcement or avoidance of the behaviour.

Rescorla and Wagner were the first to summarise the accumulated evidence about classical and operant conditioning under different contexts and stimuli combinations under the theory of Reinforcement Learning (RL, Rescorla and Wagner, 1972). The major tenet of RL is the idea that learning happens not by the co-occurrence of two stimuli, but rather due to the event being unanticipated given the current associative strengths. ‘Unanticipation’ infers there is some kind of expectation based on experience. Forming estimates and predictions about the environment allows animals to not waste

energy on unsatisfactory rewards and to invest effort in actions and objects that increase their chances of survival. According to RL, these formed predictions can be updated if there is a discrepancy or error between the expected and the obtained outcome. Hence, learning is possible only if there is an error in expectations, namely, prediction error (PE). Rescorla-Wagner rule formalises the above as in Eq. 1.1 and 1.2 (Schultz, 2017c):

$$V(t+1) = V(t) + \alpha \cdot PE(t) \quad (1.1)$$

$$PE(t) = \lambda(t) - V(t) \quad (1.2)$$

Where V is the reward prediction, α is the learning coefficient, t is trial, and λ is the reward. An important assumption of RL is that V is based on some kind of statistic that summarises the number of rewards that can be obtained in the current trial or context. Although, in a few presentations of a reward, the expected reward at trial ($t+1$) is simply the reward at trial (t), after learning based on errors, the expected V starts to regress to a mean of all possible rewards (Bush and Mosteller, 1951). As often used in economics, this can be represented as a sum of each reward's magnitude weighted by its probability, that is the expected value (EV). Although, in the Rescorla-Wagner rule, this estimate is considered over the full trial period (that is an experiment with a few probes of behaviour), Sutton and Barto within their temporal difference (TD) learning model (Sutton and Barto, 2018) considered the importance of intratrial variables that contribute to choice and learning. This update on RL accommodated sequential cues predicting a reward, an estimate of future rewards rather than just past ones, as well as a more continuous view of time represented as discrete events throughout a trial (Glimcher, 2011).

RL and TD learning theories propose that rewards that are higher than predictions will increase the associative strength between stimuli and drive learning. On the other hand, rewards that are somewhat equal to expectations will not lead to an update of estimates, whereas rewards that are worse than predicted decrease the associate CS-US strength and reduce approach (Niv et al., 2005; Niv and Montague, 2009).

Although within this PhD work, I implement brain imaging that does not involve measuring single neurones or neurotransmitters, it is imperative to consider what are the restrictions imposed by the mechanisms that give rise to the systems and responses that we observe in behaviour and larger-scale brain imaging (Craver, 2005). Looking for a signal in the brain that encodes value and value/reward prediction errors, single neural recordings beginning in the 50s established the association of the above variables with dopamine (DA) in the midbrain (Schultz, 2017a). Evidence, however, suggests that dopamine is much more versatile than simply coding rewards, now being considered also as a ‘teaching signal’ aiding the modification of behaviour when predictions are violated (Schultz et al., 1997) and also related to motivation, approach, and exertion of effort (Bissonette and Roesch, 2016; Salamone and Correa, 2012). Hence, due to the heterogeneity of dopamine and dopaminergic neurones, responding both to rewarding and aversive events (Lammel et al., 2014), it is important to always consider task design and the scale and temporal dynamics under which dopamine neurones are studied. I will first explain the anatomy of dopaminergic and DA-innervated areas and then elaborate on the findings that relate them to decision-making in animals and humans.

1.2.2 Reinforcement learning and dopamine

Since the 1960s, in a series of experiments Prof Wolfram Schultz and colleagues used single-unit recordings in monkeys and showed that dopamine neurones in the ventral tegmental area (VTA) of macaques responded to primary rewards (liquid) with a response of phasic bursts initiated about 50 – 110 ms after the onset of the stimulus and lasting around 300 ms (Schultz, 1997). More interestingly, when a reward was paired

with a preceding visual cue, dopamine neurones now stopped responding to the reward and became active at the presentation of the cue. It was also observed that an omission of a cue-predicted reward led to a dip in firing neurones, consistent with a negative prediction error signal. Later it was concluded that consistent with TD learning theory, the above dopamine responses indicate a prediction error driving learning, rather than coding reward itself (Glimcher, 2011; Montague et al., 1996). Further studies showed that dopamine neurones were also sensitive to changes in the probabilities of rewards, an effect that was shown by the dopamine response being transferred to separate cues predicting rewards with varying probabilities (Fiorillo et al., 2003). The results showed that the average dopamine response at trials with cues predicting a more likely reward was higher than that at cues predicting rewards with lower probabilities. This finding that there may be separate neurones encoding different components predicted by utility theory and necessary for a rational choice will come in handy in our later discussion on competing theories of choice under uncertainty.

1.2.3 Basal ganglia, rewards, and learning

There are a few dopaminergic pathways in the brain, but the ones associated mostly with decision-making are comprised of a few regions in the cortex and the basal ganglia, extensively summarised by Mulcahy et al. (2020) and Sesack and Grace (2010). Some of these structures are shown in Figure 1.2. The basal ganglia are subcortical nuclei composed of the ventral striatum (nucleus accumbens, ventral pallidum, and olfactory tubercle), the dorsal striatum (caudate nucleus and putamen), the globus pallidus, the subthalamic nucleus and the substantia nigra. Adjacent to the substantia nigra (SN) is the ventral tegmental area (VTA). Both produce dopamine, while the former is involved in nigrostriatal dopamine pathways (from SN to striatum), the latter forms the mesolimbic and mesocortical pathways (VTA to limbic and VTA to cortical areas, respectively).




Figure 1.2 Dopaminergic structures associated with decision-making. Substantia nigra (SN), ventral tagmental area (VTA), and other areas to which dopaminergic pathways project.

The basal ganglia (and degenerated cortico-basal connections) are implicated in several neurological disorders such as Parkinson’s disease, Huntington’s disease, Tourette’s syndrome, addiction, depression, anxiety, and obsessive-compulsive disorder (Fazl and Fleisher, 2018; Macpherson and Hikida, 2019). The basal ganglia’s function is related to motor control, learning, memory, reward processing and decision-making.

As explained earlier, prediction errors require coding of some expected value or a reference point, which is violated by received rewards, leading to realised over- or under-expectation. This ‘realisation’ results in an update of beliefs, hence, learning. Evidence shows that learning happens faster and more efficiently with rewards that violate predictions more strongly, such as extreme values or under highly volatile environments (Piray and Daw, 2021). A human fMRI study using drugs that either increased or decreased dopamine in the striatum showed the influence of dopamine on learning the estimated values of rewards, prediction errors, and its function in updating learning accordingly (Pessiglione et al., 2006). There are also two types of prediction errors,

one is the signed PE, or valence of the error, which can be better than expectations (positive PE) or worse (negative PE). Most animal studies have linked reward PEs to the ventral striatum and VTA. Unsigned PEs, on the other hand, are represented in the magnitude of the error and are proposed to function as a surprise signal, modulating and orienting attention and the motor readiness of the organism (Bissonette and Roesch, 2016). Correlates of unsigned PEs have been found in the insula, ventral striatum, and locus coeruleus, and were also associated with noradrenaline (Sara, 2009). Human fMRI studies have also shown correlations between the ventral striatum and PEs for stimuli such as faces, money, food, or olfactory rewards (Abler et al., 2006; Fouragnan et al., 2018; Hare et al., 2008; O’Doherty et al., 2003; Pagnoni et al., 2002). A recent meta-analysis showed that the evidence points towards a ‘common currency’ of social compared to non-social PEs due to a large overlap of areas, with the exception of the dmPFC which was more active in response to PEs within social experiments (Corlett et al., 2022).

The basal ganglia and dopamine are associated not only with rewards and errors but also with the initiation and urgency of motor responses. Objects that have a value usually lead to the initiation of an action, like approach or avoidance. Hence, the basal ganglia regions are thought to be involved in both reward-coding and action selection (Mink, 1996; Morris et al., 2006). Moreover, many motor structures have been associated with the valuation process of stimuli, establishing the close connection between preferences and motor activation (Kiverstein and Miller, 2015). It is also suggested that dopamine may have a dual role, coding both reward-related processes, but also the value related to an ‘incentive salience,’ or the signal that would bias a motor response towards one stimulus compared to another (McClure et al., 2003). Overall, the basal ganglia have been associated with value coding, action value coding, and prediction error coding. In the next section, we will see how these relate to risky decision-making in non-social and social contexts, how the same structures may be responsible for coding risk, and how dopamine’s functions for coding reward and uncertainty may interfere with each other to produce biases in decision-making.

1.3 Risky and social decision-making

1.3.1 Defining risk

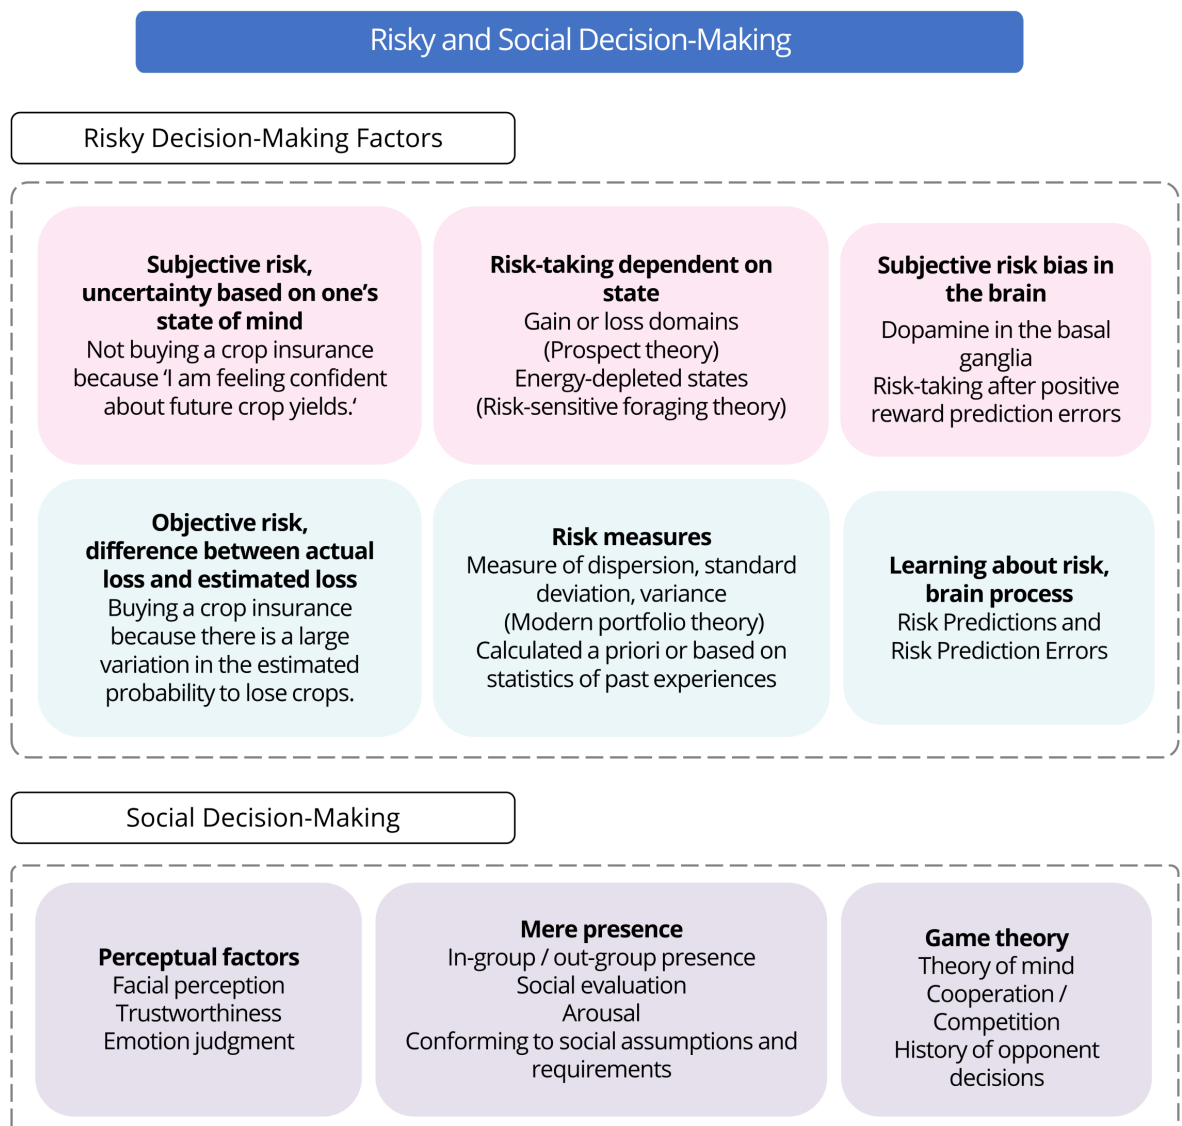


Figure 1.3 Risky and social decision-making. Risk can be subjective and objective. Subjective risk is influenced by mental states, objective risk depends on estimated values of dispersion. The science is currently exploring these states and computed values in the brain. Social factors influencing decision-making range from perception of social cues and the mere presence of others, to cooperative and competitive decisions.

There is no one single definition of risk, and there is a reason for it to mean one thing to economists and insurance companies and another to the average person. In everyday language, risk relates to some probability of loss, such as undertaking a behaviour that can lead to physical harm or losing money. More broadly risk is related to the uncertainty of an outcome. When outcomes are fluctuating risk is a measure of this fluctuation, or the relative variation of actual loss from expected loss (Rejda et al., 2005). Objective risk is specifically what is measured as such fluctuation, or what can be deduced about potential losses. Subjective risk, on the other hand, depends on a person's state of mind, beliefs, and biases (Andersen et al., 2014). For example, it is important for casinos to measure the probabilities of winning and losing with some degree of variation that allows them to predict what are possible losses that they can handle. While for the gambler, the client of the casino, it is important to believe that there is some pattern, some elusive luck, that keeps the jackpot right around the corner.

Figure 1.3 includes some factors of subjective and objective risk, as well as factors that take place in social decision-making which will come into focus later in this chapter. Subjective risk is expressed as observed choices, it is malleable, and depends on fallacies in probability estimates and beliefs, while objective risk is considered to be unalterable and based on reality (Nobanee et al., 2021). More concretely, although it is well established that cigarette smoking is a direct cause of some cancers (Hecht, 2006), hence poses a high risk for health, cognitive dissonance may lead to smokers coming up with case stories and examples of small samples of smokers who did not develop cancers in their lives. Another example of subjective risk prevails in decisions is crop insurance. It has been found that farmers are more likely to base their decisions of whether to buy an insurance on their overconfidence about the future of crops yields, which decisions were changed after they were primed with objective risk information (Fu et al., 2022). Also, people judge risks as higher if they are less frequent but have a high number of fatalities, compared to more frequent risks with low number of fatalities (Proske and Proske, 2008). One example is the fear of flying (Oakes and Bor, 2010), although it is a very low-risk mode of transport (Cobb and Primo, 2004). In such instances and many

other risky behaviours, it is important to understand subjective risk perceptions and how to steer them towards more objective estimates of risk. Next, I will outline some existing measures of objective risk and explain the theories that try to explain risky behaviour in humans, touching upon brain function.

The need to clearly define risk and estimate it from possible outcomes (wins and losses) arose with the development of financial theories that aimed to maximise the portfolios of investors. The conceptualisation of risk as variance or the standard deviation of a distribution of variable rewards arose from normative economic theories, those that prescribe what would be the most profitable or objectively desirable choice, given the situation at hand. Harry Markowitz recognised that uncertainty can be measured by the variability of outcomes in his modern portfolio theory, incorporating the mean-variance approach (Markowitz, 1952; Steinbach, 2001). Ever since, this approach has been used to create diversified portfolios, maximising returns to investors, and aiding financial decisions depending on risk propensities. It was found that teaching the method of mean-variance was more beneficial in improving decisions about portfolios compared to reinforcement learning (Olschewski et al., 2021).

Nevertheless, under naturalistic circumstances, it was found that humans and animals do not necessarily make choices that maximise expected values (Schuck-Paim et al., 2004; Stauffer et al., 2015). A few issues arise when translating animal models into research with humans. When it comes to animal studies, rats, monkeys, and other animals can be (compared to humans) logical in their choices, predictable and consistent across tasks. This observation has been attributed to extensive training of animals before their behaviour and neural responses are recorded (Schultz et al., 2021). Moreover, animals are tested with what can be considered primary rewards such as sucrose, food, or even stimulating drugs, and they are usually under energy-depleted states with the aim to increase the value of even small amounts of rewards. The energy states of animals have also been found to influence their rationality and compliance with economic models (Schuck-Paim et al., 2004). On the other hand, humans are less constrained by

the survival needs that animals experience and have more personality characteristics that drive even more sophisticated preferences. This subjective interpretation of similar rewards that vary across individuals and situations led economists to develop the concept of utility, or the subjective value of goods (Małecka, 2020). As it recognises the subjective value of rewards and provides a clear mathematical model of expectations, utility theory has been essential in modelling the activity of dopamine neurones in response to prediction errors in both animals and humans (Caplin and Glimcher, 2013; Ferrari-Toniolo and Schultz, 2021; Phillips et al., 2007; Stauffer et al., 2014).

The above led many to reconsider the idea of a ‘rational human econ’ that makes their decision by carefully calculating all possibilities, averages, and known variability of the environment. It became more common to try to understand decision-making by incorporating factors like impulsivity, affective states, emotions, and heuristics that substitute carefully weighted costs and benefits (Finucane et al., 2000; Loewenstein et al., 2001). These observations in decision-making led to the hypothesis of a dual system by Kahneman (2012) and the somatic marker hypothesis (Bechara and Damasio, 2005). Taking the affective response to loss aversion, prospect theory increased the predictive power of economic models remarkably. According to the theory, humans who are generally risk-averse may become risk-seeking in the context of potential losses (Kahneman and Tversky, 2018). This effect has also been found in decision-making in animals and has proven powerful in predicting decisions (Trepel et al., 2005). Gains and losses were also found to have opposing effects on the BOLD activity of some brain regions. An fMRI study found that most regions that coded gains had decreased activity as the size of losses increased (Tom et al., 2007). These regions included the striatum, the ventral ACC, the vmPFC, and medial OFC, which are also associated with reward anticipation and valuation. I will focus on the overlap between reward and risk processing regions later in this chapter, however, the more significant issue with

drawing conclusions about risk here is that under prospect theory the propensity for risk is only a consequence of the context of rewards and their probabilities rather than a separate estimation of risks (Bossaerts, 2010). Nevertheless, evidence suggests that at least under some circumstances risk is coded independent of rewards.

This brings us to another view of the computations underpinning decision-making under uncertainty. Rather than using utility functions as in the case of utility and prospect theories, some propose that the brain computes the moments of a reward distribution separately from one another: expected value (average), variance, and skewness. This view emphasises that in order to choose between risks, the brain estimates risk and then incorporates these estimates into a prediction error algorithm related to Rescorla-Wagner learning models, such as TD learning, and chooses risky options according to current goals (Preuschoff and Bossaerts, 2007). Going back to the mean-variance approach, a few studies have incorporated EVs and variable SDs in their experimental paradigms to study the representation of these statistics in the brain.

Due to its mathematical simplicity, and high optimisation of choice, the mean-variance approach from portfolio theory (Markowitz, 1952) has been used extensively in animal and human studies of decision-making under uncertainty and was found superior to utility theory when it comes to risk learning (d’Acremont and Bossaerts, 2008). Experiments that keep rewards constant and vary the standard deviation of different stimuli show that risk can be encoded independently from rewards. Christopher Fiorillo and his colleagues (Fiorillo et al., 2003, 2005; Fiorillo, 2011) studied the response of dopamine neurones in relation to specifically variable rewards compared to safer rewards with a similar expected value. The studies showed that dopamine responds to risky cues more strongly compared to safe cues. Not only were animals behaviourally drawn to different groups of rewards that varied by risk, but their dopaminergic neurones in the orbitofrontal and posterior cingulate cortex were differentially active in a parametric fashion for different levels of risk (Fiorillo et al., 2005; Lak et al., 2014). The interpretation that dopamine neurones code uncertainty has been questioned due to

the asymmetrical dopaminergic coding of positive and negative PEs, and issues with averaging across trials (Niv et al., 2005). Additionally, serotonin has also been extensively linked to uncertainty coding and recently has increased as a target of research in uncertainty (Grossman et al., 2022). Hence, although I focus on theories of dopamine, it is likely that other neurotransmitters interact with dopamine to code risk and aid orienting responses.

1.3.2 Human studies with mean-variance dissociation

Animal studies usually benefit from extensive training (that can take days) and energy depletion through hunger and thirst which in turn leads to high motivation and fast learning. This has also led some to believe that monkeys are more rational and smarter than humans because they satisfy basic rationality requirements when their energy is depleted (explained in 'Not smarter, just more constrained' by Schultz et al., 2021). It is also possible to implement some form of re-evaluation, where the preference for different rewards is controlled and manipulated. This strategy is more difficult to implement with humans because rewards of high value are rare (at least in the lab) and depleting energy through starvation is not feasible. Naturally, most experiments use monetary incentives, however, as I will elaborate on later, very small rewards can lead to low ecological validity and unexpected behaviours. Another issue that arises from the use of low-value rewards is the depletion of attention and motivation, which in turn can deplete cognitive resources more easily.

To make learning simple, and avoid depletion of cognitive resources, experiments trying to test the mean-variance approach have implemented an association of one stimulus with either of two values with equal probabilities. If a coin is tossed, the probability of each of its outcomes is 50%. However, assigning different rewards to each side of the coin can result in more risky or less risky environments (O'Neill and Schultz, 2015). For example, in one case we can assign £40 to heads, and £60 to tails, and in another instance, heads can be £20 while tails be £80. Now the average of both cases is £50,

however, the distance between the two possible outcomes is greater in the £20-£80 option. If the probability of a reward varies, the risk is greatest when the probability of a reward is 0.5, whereas it decreases with a probability of a reward being 0 or 1 (Figure 1.4 - right). This is consistent with the conceptualisation of economic risk representing uncertainty.




Figure 1.4 Risk and reward relationship. Left - Reward probability distribution with the expected value (EV) and variance. Right - Risk as a function of reward probability (figures from Schultz, 2015). Risk is highest when rewards probability is at 0.5, while it is lowest when reward probability is 0 or 1.

If processing risk is essential for making better decisions and we know that risk is coded in the brain, the next step is to explore which components of risk representation are informative for learning and decision-making and which are different or similar to reward coding. If there is risk processing similar to reward representation, one would expect that there will also be an error, the difference between expected and experienced risk, or risk prediction error (RiPE). Going back to the example of insurances, for an insurance provider it would be essential to measure the difference between expected loss and actual loss, and especially what is the accuracy of this measure. For example, if out of 10 000 houses about 100 burn every year, it is expected that one year there will be 90 houses burning, and the following maybe 110 will burn (Rejda et al., 2005). If the next year 95 houses burn, the risk prediction error is small. However, a certain year 200 houses burn, then the risk prediction error is large and estimates have to be adjusted.

Studies have tested risk and RiPEs in humans with fMRI and purely probabilistic designs, using cards, numbers, and symbols with differing probabilities. Preuschoff et al. (2008) used a card game in the fMRI to expand on the animal literature on risk and RiPEs. During the task, two cards were presented in sequence. Participants had to guess if the second card would be higher or lower than the first card before seeing any of the cards. In this way, the experiment achieved the ability to measure the evolution of average expected reward, predicted risk, and deviation from the two values at two time points (at the presentation of the first and then at the second card). The results revealed a correlation between activation in the anterior insula and the model of risk, and a different pattern of insula activation and the model of risk prediction error. Moreover, insular activation resembled the modelled quadratic function of risk relative to reward probabilities (shown in Figure 1.4 - right).

Another fMRI study used a version of the Iowa Gambling Task with different expected rewards and variances across 4 decks of cards (d’Acremont et al., 2009). Consistent with previous results, a differential response was found for reward prediction errors (in the striatum) and risk prediction errors (in the insula and inferior frontal gyrus).

Most recently, Lauffs et al. (2020) tested the 2-card sequential game in an EEG and pupillometry study and found a two-component response following card presentations. The first component was interpreted as the salience of the stimuli (P1*, ~ 180 ms), while the second one (P3, ~ 300 ms) was inferred to reflect the risk prediction error. Intriguingly, the first salience component was not simply generated in visual areas, but also in temporal cortices, suggesting insula activation. The pupil activity correlated with the two-component EEG response, suggesting an underlying noradrenergic activity (Reimer et al., 2016). This study laid the foundations for the temporal decoding of risk prediction error in the human brain and suggested that risk and RiPEs serve a ‘surprise’ function during uncertainty coding.

In a more naturalistic setting it was shown that human behaviour is also influenced by the deviations of expected and experienced risks (RiPEs). Vavra et al. (2018) showed that in the ultimatum game, the decision to reject an unfair offer was influenced differentially by the mean and the variance of the distribution of offered rewards. The behavioural results showed that the mean of offered rewards affected the threshold at which rewards were accepted while the variance affected the strictness with which these rewards were accepted. Expectations of greater variation in offers increased the likelihood of accepting offers, and especially lower ones. Offers with lower variance were accepted less often. Hence, decisions about whether to play or not were influenced by the uncertainty in the outcome and preceding expectations. This study did not only show that RiPEs could influence decisions, but they did so in a social task where rewards were dependent on another person.

The above studies, using the mean-variance approach, show that risk predictions and errors are computed separately in the brain with both temporal and spatial differences. Additionally, risk prediction errors within a social game were shown to influence decision-making. Currently there are no neuroscience studies that have tested risk processes, such as predicted risk and errors, within social contexts. Next, I will turn to outline theories of increased trial-by-trial risk-taking and later I will discuss the literature that currently aims to clarify whether there is one unifying system for risky decision-making in general, or whether there is a separation of networks for social and non-social choices under uncertainty.

1.3.3 The basal ganglia model of risk-taking

The structures within the basal ganglia show activity that aligns with the principles of reinforcement learning models, such as coding of prediction errors via midbrain dopamine neurones (Morris et al., 2006; Schultz, 1997). Given an operant conditioning environment, when an organism performs different actions and is rewarded differentially for them, it is considered that separate values are assigned to each action according to

the associated rewards (Wunderlich et al., 2009). For both reward values and action values, the substantia nigra transmits signals to the striatum associated with positive and negative PEs (Reynolds et al., 2001). These associations signalled from the SN are then coded by the ventral striatum (nucleus accumbens and olfactory tubercle) and the dorsal striatum consisting of the caudate nucleus and putamen (Lau and Glimcher, 2008; Morris et al., 2006; Samejima et al., 2005). And the link between the dorsal striatum and action values was found to be causal (rather than only correlational) via the use of optogenetic stimulation (Tai et al., 2012).

Figure 1.5 shows a schematic representation of the long-term potentiation connections (blue) between the substantia nigra (SN) and the striatum. According to the basal ganglia learning model, approach and avoidance depend on learning action values, coded by dopamine within two pathways: Go and No-Go (Bogacz, 2017). After learning of stimulus-outcome contingencies is complete, dopamine plasticity between the SN and the Go pathway in the striatum is increased by positive reward PEs (green), while negative reward PEs activate dopaminergic neurones within a No-Go pathway (red). The SN-striatal pathway exerts a net facilitatory effect on the thalamo-cortical connections, enhancing or inhibiting behaviour (Gilbertson and Steele, 2021). Acting on the thalamus and motor areas, positive RPEs and increased Go signalling aid approaching behaviours, while negative RPEs and enhanced No-Go connections lead to a lack of approach (inhibition).

Fluctuating levels of dopamine in the basal ganglia are crucial for exploration and exploitation switching in healthy behaviour under uncertainty (ibid.). However, increased dopamine levels within the striatum have been related to the severity of gambling disorders (Joutsa et al., 2012) and up to 13% of people treated for Parkinson's disease with dopamine agonists experience pathological gambling (Dodd et al., 2005; Driver-Dunckley et al., 2003; Santangelo et al., 2013; Zhang et al., 2021). A recent study with healthy volunteers found increased risk-taking with a tonic increase of dopamine with levodopa compared to placebo, although the drug effect depended as well on baseline

dopamine levels (Hirschbichler et al., 2022). Moreover, the Go and No-Go pathways are populated by D1 and D2 dopamine receptors, respectively, and different variants of genes that express as higher D1 or D2 receptors lead to better Go or No-Go learning (Frank et al., 2007). These studies point to the notable importance of the basal ganglia and tonic dopamine levels in trait and state risk-taking.

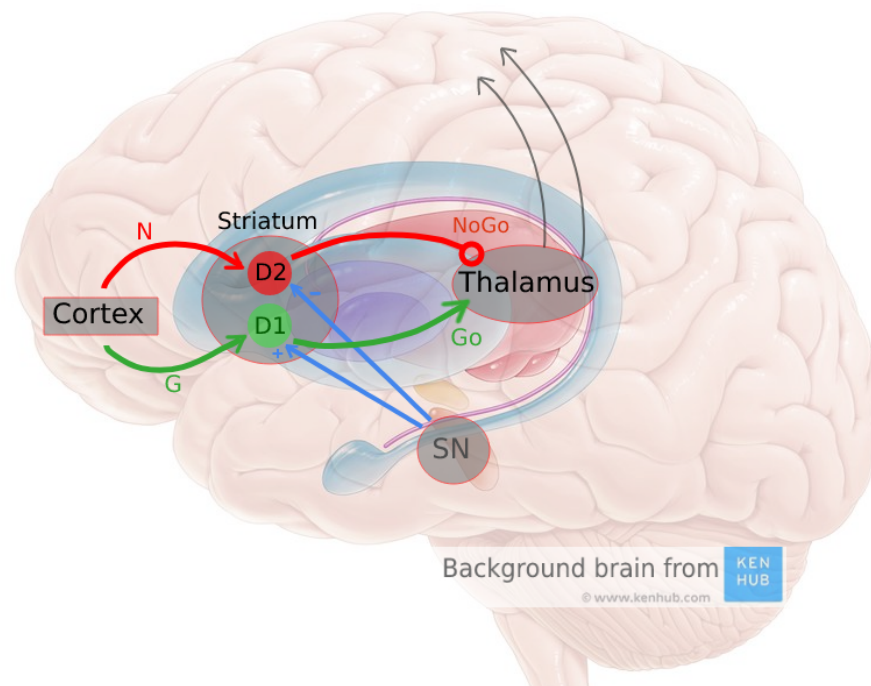


Figure 1.5 A schematic representation of long-term potentiation and depression. Dopaminergic pathways from the substantia nigra to the striatum, respectively acting on the Go (green) and No-Go (red) pathways, leading to approach or avoidance. Blue arrows indicate long-term potentiation, or learned positive and negative outcomes through positive and negative PEs.

The striatum is influenced by connections from the SN (D1 and D2 pathways) and VTA, but also from areas like the OFC and vmPFC (Ferry et al., 2000; Philiastides et al., 2010a). This influence from the cortex on the striatum has been shown to modulate the way PEs are normalised (Park et al., 2012). In other words, cortico-striatal connections were shown to be crucial for the adaptive coding of moment-to-moment rewards with the OFC being sensitive to the contextual distribution of rewards (Padoa-Schioppa, 2009). The activation of the medial PFC in rats was associated with changing environments that require higher cognitive demand and signalling appropriate behaviour via

projections to the ventral striatum (Howland et al., 2022). This is consistent with fMRI studies suggesting that the dmPFC is related to action selection and goal-directed decisions rather than initial decision variables (Larsen and O’Doherty, 2014). Another area with connection with the striatum and associated with information-seeking behaviour when uncertainty is present is the anterior cingulate cortex, ACC (Monosov, 2020). The ACC has been shown to code PEs and to track the volatility of the environment, and its projections to the ventral striatum have been suggested to influence the estimate of rewards and the rate of learning (Behrens et al., 2007). The OFC, dmPFC, and the ACC, are only some of the prefrontal areas innervating the striatum and are thought to have a top-down influence on decision variables and action selection. Within the second reported experiment, I aimed to test the theory, compare competing models, and provide a spatiotemporal outline of stimulus-PEs with EEG. Another influence on cortico-striatal connectivity may be a social context which we will explore next.

1.3.4 Social decision-making, risk, and learning

Observing our everyday lives, it does not take much to consider the truthfulness in the saying by Aristotle that humans are social animals. The social brain hypothesis proposes that primates’ unproportionally large brain-to-body size ratio is the result of a demand from complex social interactions (Dunbar, 1998). And the social world is indeed an inseparable part of human existence. We use others’ behaviours and their opinions as reference points when making decisions (Parks and Sanna, 2018; Wood and Hayes, 2012), we possess the ability to mirror other’s motor actions (Iacoboni, 2009), and we are influenced by others’ feelings of pain, joy, surprise, and anger (Decety, 2010).

Social cognitive neuroscience arose as a field to combine social psychology and neuroscience, aiming to reveal the brain processes related to social interactions from perception to decision-making and determining whether these processes are the same or different from non-social contexts (Ochsner and Lieberman, 2001). Most research on social decision-making started from the realisation that social factors do matter for

economic decisions, leading to the development of ‘game theory’ (Sanfey, 2007). Game theory aims to understand the mechanisms that take place when people make decisions in the presence of others. This can range from the mere presence of another person to games of cooperation and competition, but they all entail one’s need to understand other people in order to gather relevant information for improved decision-making. Understanding others may start with the most basic cue of their current state which is the facial expression, or it can be informed by more complex information like what was their most recent decision. Our ability to infer the intentions and actions of others is termed Theory of Mind (ToM, Premack and Woodruff, 1978) and it is an essential component of Game theory as it explains the ability to make decisions acknowledging the state of the other ‘player.’ To be able to ‘mentalise’ another’s mental state, to put oneself in another’s shoes, requires the capacity to imagine someone else’s goals, motivations, and affective states. Under more complex situations we use our knowledge of social norms, or what would be an appropriate action for most people in the current culture. We also need to consider that another person may be feeling in a different way and behave unpredictably depending on their current state. This is when we refer to the face for cues.

We often underestimate how important facial features are in social interactions because they are such an integral part of our lives. In an experiment with a trust game scenario participants were found to rely more on facial cues, rather than previous payoffs, and were faster and more accurate in their decisions when presented with faces compared to simply numerical payoffs (Jaeger et al., 2019). Although the participants reported that they relied more on the numerical payoffs, the data showed otherwise. Babies as young as 6 months show a preference for faces, indicating discriminability between faces and other objects (Ludemann and Nelson, 1988). In a study with adults, faces were shown for 50 ms, 100 ms or 500 ms, and were then masked. During masking, the participants answered ‘yes/no’ on whether the face was trustworthy and then they indicated their confidence in the choice. It was shown that adults were able to detect facial trustworthiness even when faces were presented for only 100 milliseconds, the

equivalent of 'a single glance at a face' (Todorov et al., 2009). Others have shown that the trustworthiness ratings of the faces of presidential candidates can predict voting outcomes (Ballew and Todorov, 2007; Sussman et al., 2013); and un-trustworthiness of faces can influence legal decisions and lead to faster guilty verdicts (Porter et al., 2010). These results show that, although underappreciated, faces are a big factor that can bias social decision-making.

Risky decisions during games of poker are also influenced by different facial expressions. In a study with online poker, the faces of the opponents were presented to participants who previously rated the face images on a trustworthiness scale (Schlicht et al., 2010). Although threatening faces did not lead to altered decisions compared to baseline, participants made more mistakes when presented with faces with positive expressions. These results show that although opponents were not physically present during the game, simple cues such as face images could influence risky decisions. These findings point out that facial expressions and their perceived trustworthiness can be considered cues that reduce uncertainty. The study by Jaeger et al., 2019 showing that participants relied on facial cues rather than numerical payoffs, but did not report so, suggests that social 'shortcuts' to decisions may be automatic and implicit, which is consistent with previous findings (Rezlescu et al., 2012). Similarly to non-social decision-making, a dual system was proposed that includes automatic and controlled processes within social contexts (Bargh et al., 2012; Lieberman, 2007). For example, the finding that the amygdala (usually responding to emotional content) may be activated differentially by fearful compared to happy facial expressions, without clear awareness about the stimuli suggests a fast automatic process (Williams et al., 2004).

Another study integrated social and utilitarian factors in a risky decision-making task in which the participants were presented with pairs of faces with different probabilities of rewards. The results showed that when participants selected a happy face rather than a sad or angry face, they were biased and put more weight on the positive outcomes, while they put less weight on the positive outcomes following angry or sad faces (Averbeck

and Duchaine, 2009). These results show that faces are used as additional, intuitive cues in valuation processes. The next question is how facial cues get integrated into resolving uncertainty when making decisions. Are faces coded as value-inducing cues like non-social objects, and additionally can faces reduce uncertainty if they happen to be accurate cues of rewards?

The perception, evaluation, and interpretation of faces is the first point of social judgment and resolving uncertainty, hence, using faces as stimuli can be the ground for bridging social and non-social neural and computational processes in response to risk. Through empirical evidence, two main schemas have been developed to explain the neural underpinnings of social versus non-social decision-making in general, one proposing a ‘common currency’ and the other a ‘social-specific cognition’ (Ruff and Fehr, 2014). The first one (the extended common currency schema) proposes that there is a general mechanism which interprets and assigns value to both social and non-social factors, even though perceptual information may come from separate networks. The second one (social-valuation specific schema) proposes that the computational principles of social and non-social factors and valuation are similar in structure but occur in separate networks that evolved and specialised with the demands of a social environment.

Although there are clearly specialised areas for social stimuli (like the fusiform face area and mirror neurones), currently, there is an overarching consensus that social rewards and risks are coded in identical networks with non-social ones (Rilling and Sanfey, 2011; Ruff and Fehr, 2014). It is crucial to bear in mind that social specificity may depend on the different levels of description (Lockwood et al., 2020). For example, social decisions 1) may pose notably different goals than other types of decisions, 2) may use distinct algorithms from non-social situations, and 3) may differ in their neural implementation. The scope of this work does not cover the differentiation between the levels, however, within the reported experiments I aimed to add the lowest level of influence from social components to build on previous non-social findings or to

maximize our ability to compare social to non-social processes. I did this in such a way that I would reduce the influence of implementation of different goals and strategies under social conditions (in experiment 1), and if indeed different algorithms were used during non-social compared to social contexts, I could capture this in behavioural and neural levels (as in experiment 2).

Prediction errors are one example where the two domains overlap. A meta-analysis showed that areas including the dorsal and ventral striatum, pallidum, vIPFC, orbitofrontal cortex, and insula appear to overlap in social and non-social tasks testing PEs (Corlett et al., 2022). Nevertheless, more activation was found in the dmPFC during social tasks, which region is also related to cognitive (as opposed to affective) theory of mind (Abuz-Akel and Shamay-Tsoory, 2011), and was found to be selective for predictions about other's beliefs and own reward prospects in social contexts (Elliott and Deakin, 2005; Jamali et al., 2021). Overlaps were also confirmed by an EEG study, showing that violation of social norms results in altered feedback-related negativity (FRN) in EEG, which was consistently associated with non-social PEs, suggesting that not only spatially but also temporally social prediction errors may be processed like non-social ones (Harris and Fiske, 2010).

As discussed earlier, studies have localised risk processing in both animals and humans to the striatum, insula, and lateral orbitofrontal cortex (Rangel et al., 2008). Volatility and prediction errors have also been associated with the ACC (Behrens et al., 2007). It was proposed that the ACC is related to tasks that require online risk learning, while the anterior insula and inferior frontal gyrus/lateral orbitofrontal cortices – during risk judgments of known probabilities of rewards (Fitzgerald et al., 2010). However, this dissociation is still to be confirmed with fMRI. The insula is active for social cues, reward-related signals, and perception of internal physiological states (Rogers-Carter and Christianson, 2019). It was associated with predictions of risk and errors (RiPEs) in probabilistic tasks (Clark et al., 2008; Loued-Khenissi et al., 2020; Preuschoff et al., 2008), and the integration of social sensory information, empathy and emotion

recognition (Bernhardt and Singer, 2012; Terasawa et al., 2021). However, it's worth noting that these functions may not be exclusive to the insula and could also involve the mid-ACC (Corradi-Dell'Acqua et al., 2016). The converging activation for social, rewarding, and risky cues within the same regions suggests that there may be domain-general processing, although no studies have examined the combined effect of risk processing (including predicted risks and risk prediction errors) with social stimuli. With my first experiment, I aimed to reveal the spatial and temporal profile of predicted risk and RiPEs with EEG, while implementing social stimuli in the form of faces of varying subjective trustworthiness. I compared the results to previous studies with non-social stimuli.

Although action values based solely on expected rewards and risks (coded by dopamine in the striatum) can be used to make decisions and can be used by scientists to predict decisions, other factors such as social biases, exploration strategies, and optimistic thinking, all processed through prefrontal and anterior cingulate cortices may influence typically 'rational' decision-making (Rushworth and Behrens, 2008). Within social contexts with the presence of faces of varying trustworthiness, there may be better estimates of predicted risk, hence, lower risk prediction errors. For example, if a face is perceived as highly trustworthy, the probability of making an error with it would be lower than if the confidence in its trustworthiness is low, like in the case of more neutral faces. We can expect there to be an additive effect of numerical outcomes and social cues on lowering risk prediction errors. For example, more trustworthy faces suggesting higher rewards and less trustworthy faces suggesting lower rewards, compared to more neutral faces carrying a small difference between the rewards they lead to. Faces can be a shortcut in decision-making that leads to the feeling (or more automatic and unconscious processes) of 'resolved uncertainty' and can be a supplement to numerical cues, which I explore in my first experiment. The second focus of the current PhD was to test a model that aims to explain risky decision-making as stemming from exposure to positive prediction errors. I introduced an additional component of a group membership bias aiming to test its additional influence on risk-taking.

1.3.5 Social group membership, cortico-striatal inputs, and risk-seeking

Another powerful factor in decision-making in social circumstances, apart from initial superficial cues about others, is the relationships of those we are surrounded by (Rilling and Sanfey, 2011). In-group members are individuals who belong to the same group as us, such as our family, friends, ethnic or cultural group, or any other group which we identify with. Out-group members, on the other hand, are individuals who belong to a different group than ours. Group membership bias, also known as in-group bias, refers to the tendency of individuals to favour members of their own group (in-group) over members of other groups (out-group, Myers, 2016). This bias can influence our decisions and behaviours in subtle and sometimes unconscious ways.

Depending on the strength of the relationship between ourselves and groups, biases may be translated into attitudes and behaviour. For example, there may be biased judgments, such as giving more weight to the opinions or contributions of in-group members (Ben-Ner et al., 2009a; Jannati et al., 2016) and being more lenient towards their mistakes or shortcomings (Molenberghs et al., 2013). In terms of explicit behaviour, we tend to trust and cooperate more with in-group members (Fujino et al., 2020; Plötner et al., 2015), as we perceive them as more similar to ourselves and therefore more dependable. Additionally, we may be more likely to seek opinions and input from in-group members and may be more influenced by their perspectives (Mackie et al., 1990), leading to a biased decision-making process that may not consider the full range of options or viewpoints.

Also, the degree of social influence is intricately linked with the need to feel part of a group and to be approved, as is prominently revealed during adolescence. I will use adolescence as an example as it is a period in which social influence, risk-taking, and the maturation of cortico-striatal connections converge in a meaningful relationship.

It is well-known that risk-taking increases in the transition between childhood and adolescence and declines with the emergence of adulthood (Defoe et al., 2015). Compared to adulthood and childhood, adolescents engage more in life-threatening behaviours like alcohol and substance abuse, having unprotected sex, and driving recklessly or under the influence of substances (Steinberg, 2008). A major reorganisation of cortico-striatal areas during this period of life has been shown to be intricately linked with risk-taking and impulsivity (Decker et al., 2016). From a neurocognitive perspective, the increase in risky behaviour has been explained by the ‘dual systems model’ (also ‘maturational imbalance theory’). This model posits that an increase in sensation-seeking, modulated by the maturation of the socio-emotional system (ventral striatum and medial prefrontal cortex), and a decrease of cognitive control by delayed maturation of cognitive control systems (dorsal prefrontal, parietal and anterior cingulate cortices), results in the observed behavioural change (Balter and Tamis-LeMonda, 2016).

In some cases when risk evaluation and self-monitoring are comparable to those of adults, the emotional reward from risk-taking may be high enough to override rational decision-making and behaviour (Gardner and Steinberg, 2005). Therefore, Steinberg (2008) suggests, this is the reason why simply educating adolescents about the costs of risky behaviour has been inefficient so far. Even if the risk is correctly evaluated and considered, risk-taking may be observed due to other factors (Dahl, 2004; Defoe et al., 2015; Reyna and Farley, 2006). One of those factors for adolescents may be the presence of peers (Albert et al., 2013).

The social identity model of risk-taking proposes that shared group membership (being surrounded by in-group members) attenuates risk perceptions, decreasing feelings of uncertainty, and consequently increasing risk-taking (Cruwys et al., 2021). This was timely shown and supported in the context of the pandemic. People were more likely to rate activities such as sharing a drink, leaving a used tissue, and shaking hands with someone as less risky when performed by in-group compared to out-group members (Cruwys et al., 2020a). In experimental studies, where norms are not directly verbalized

or there is reduced knowledge about the ‘in-group,’ behaviour may be influenced by just the presence of others. Mere presence (audience effect) has been long known to influence behaviour, especially in groups whose sensitivity to peer evaluation is increased (Weigard et al., 2014; Wolf et al., 2015). If the only mediator is fear of social evaluation, this means that in private setting people are going to show behaviour more like the one before the specific social norms were introduced. Another suggestion is that mere presence only leads to arousal and increased activation of emotional centres, which in turn increases risk-taking, emotion-based decision-making and behaviour. In line with this idea, neural change has been shown to occur in the presence of peers even when no decision-making takes place (Somerville et al., 2013). With my last experiment, I further aimed to test the proposed additive influence of in-group and out-group presence on risk-taking hypothesising that in-group presence would increase risk-taking on top of positive stimulus-PEs.

1.4 The goals of this thesis

The previous sections summarised the theoretical background of value-based decision-making and its association with dopamine and the basal ganglia. I focused on the challenges and solutions to define risk and to develop the best model to capture risky choices. And I reviewed studies that probed the brain for the correlates of risk-taking in social and non-social tasks, showing the involvement of specific areas within the prefrontal cortex and midbrain. Additionally, I introduced the involvement of social influence on risk-taking and the importance of studying individuals through the lens of a social world.

Within this thesis, I introduce two experiments divided into Chapters 2, 3, and 4. The primary goal of this work was to test risky decision-making by combining recent paradigms and models of risk, compiled in an interdisciplinary manner from economics, computational modelling, and animal work. I aimed to answer unresolved questions about:

1. Whether risk and deviations from predicted risk (errors) are computed similarly to rewards when a social component is introduced. And to disentangle these computations by EEG.
2. If a recent model on risk-seeking driven by transient fluctuations of dopamine can predict risky choices better than other established models from economics and reinforcement learning. And if this prediction-error-influenced risk-seeking can be predicted by single trial EEG fluctuations.
3. How the presence of in-group and out-group members affects uncertainty and risk-seeking on top of the proposed effect of prediction errors. And do the neural patterns observed during baseline compare to those recorded during social observation?

Ultimately, these projects make valuable contributions to the existing literature on human decision-making under risk. Through a comprehensive integration of recent methodologies from animal studies, functional magnetic resonance imaging (fMRI), and computational research, I present compelling evidence on how risk-related variables manifest in electroencephalogram (EEG) measurements. Moreover, I address critical inquiries regarding the neural and computational foundations of social neuroscience. The first experiment confirms the presence of a brain process that codes risk and a separate one that codes prediction errors related to risk in humans. Using social cues mapped to numerical payoff and finding similar results to studies with non-social cues, I supported the notion that both social and non-social rewards are encoded similarly in the brain. Building upon an intriguing new theory of the basal ganglia, the second experiment successfully replicates recent findings and shows computationally that risk-taking may be the consequence of positive reward prediction errors. Significantly, I

uncover the EEG correlates of stimulus-locked prediction errors that potentially drive individuals to seek risks under some contexts. Conversely, the latter part of the second experiment raises important inquiries about the impact of group membership bias on risky decision-making. By exploring mental health factors, these final findings hold substantial significance within the context of ongoing efforts to comprehend learning, gambling, addictions, and other conditions in which any stage of the four-module decision-making process may be compromised.

Spatiotemporal Representation of Risk and Risk Prediction Errors with Social Stimuli

This chapter outlines the background, findings, and discussion of our first experiment aiming to test risk processing in the brain. Specifically, we aimed to test whether the risk can be decoded separately from rewards and signed and unsigned risk prediction errors. (Within the following chapters I will continue by using the plural first-person pronouns as these experiments are the combined hard work of myself and my supervisor.) We incorporated social stimuli in the form of faces that the participants rated prior to the main task, and which were mapped onto numerical reward values. We further used a design that was able to decouple risk from signed and unsigned risk prediction errors and we found that there were separable EEG responses for the three risk-related variables. These results are in line with previous fMRI and EEG findings and add additional evidence to the theories of a common currency of social and non-social decision-making.

2.1 Background

Imagine that you are visiting a new city with a few different options to buy coffee. One approach would be to stick with the high-street chains, to minimize the uncertainty regarding coffee quality. Another approach would be to try a local boutique coffee shop to improve your experience, but this move might also lead to possible disappointment due to higher uncertainty in the shop's bean quality. Living in complex ever-changing environments, we are constantly faced with uncertainty when making choices. Although a sensible strategy would be to aim for the highest-valued option, whether it is coffee, money, or pleasurable experiences, the outcome of our choices is not always fully predictable. In the coffee shop example, it is not easy to predict if we will feel better from a novel experience or from sticking with the familiar. The process of optimising our decisions is believed to rely on a reinforcement-guided framework, whereby our expectations are updated every time we encounter a new experience (Rescorla and Wagner, 1972; Sutton and Barto, 2018; Watabe-Uchida et al., 2017). Reward prediction errors (RPEs) are formalised as the difference between experienced and expected rewards and they serve belief formation and the ongoing comparison of beliefs against experiences (Niv and Montague, 2009; Watabe-Uchida et al., 2017). The valence (sign) of RPEs leads to either repeated behaviour (positive RPEs) or an avoidance response (negative RPEs), while the magnitude of RPEs (unsigned RPEs) controls the speed of learning (Ouden et al., 2012; Schultz, 2016a).

Most everyday choices, however, are not easy to weigh against previous experiences as they pose varying levels of uncertainty. One may go for artisan coffee if one is interested in 'gambling' with taste as the outcome is more variable. So, in this case, the choice is not one between known outcomes but rather based on the tolerable degree of variation around an average reward, also known as 'risk' (Bossaerts, 2010). In statistical terms, risk can be represented as the standard deviation or the variance of a distribution of

rewards (Steinbach, 2001). Several studies have used this definition and have shown that under uncertainty, decisions can be informed by risk, independent of rewards (d’Acremont and Bossaerts, 2008; Fiorillo et al., 2003; Kacelnik and Bateson, 1996; Niv et al., 2012; Tobler et al., 2009).

It follows that learning about and updating our representations of the uncertainty of rewards can add a behavioural advantage in novel or volatile environments. Just like errors in estimating expected rewards, there can also be errors in the estimated variation around these rewards. Similar to RPEs, risk prediction errors (RiPEs) are defined as the mismatch between experienced and expected risk and can be used to update our future risk estimates. RiPEs can be signed or unsigned with each component of this signal (valence or magnitude) serving a potentially separate role in reducing uncertainty in future choices (d’Acremont et al., 2009; Preuschoff et al., 2008).

Human fMRI studies, using purely probabilistic (non-social) decision-making tasks, found correlates of expected risk and risk learning in the anterior cingulate cortex (ACC), the anterior insula and the inferior frontal gyrus/lateral orbitofrontal cortices (Behrens et al., 2007; Brown and Braver, 2007; Critchley et al., 2001; Fitzgerald et al., 2010; Fukunaga et al., 2018; Payzan-LeNestour et al., 2013; Preuschoff et al., 2008). Similarly, neural correlates of RiPE were found in the ventral striatum, the inferior frontal gyrus (IFG), and the insula (d’Acremont et al., 2009; Preuschoff et al., 2008). While fMRI studies enabled the identification of specific brain nodes encoding different risk signals, they are limited in their ability to explain how these signals unfold in time.

A recent EEG study was the first to directly investigate the period of risk processing signals in a purely probabilistic (non-social) scenario (Lauffs et al., 2020) and found a P3 event-related potential component over central electrode sites covarying with the magnitude of RiPE signals. In this study, RiPE signals were derived from the errors in reward prediction on each trial, which poses the possibility that the measured signals are instead a consequence of valuation (O’Neill and Schultz, 2010a). To address

this caveat O’Neill and Schultz (2013) designed a task in which the expected reward of the different options remained constant, while the standard deviation around this expected reward (i.e., risk) was varied systematically across trials. RiPE signals were then computed as the difference between the trial-wise risk and the overall risk across all trials (expected risk). While this design enabled the identification of neurones in the orbitofrontal cortex responding to unsigned RiPEs, it could not reveal signed RiPE signatures, since they were confounded by the trial-wise estimates of risk (due to a fixed value for overall predicted risk).

Here, we use a modified version of the task by O’Neill and Schultz (*ibid.*) designed to properly decouple risk, signed and unsigned RiPE signals while doing so using socially relevant stimuli (framed in the context of a trust game), rather than non-social gambles as in the studies described above. We aimed to study the extent to which risk and RiPE signals also emerge in the social domain and offer an account of how these signals unfold in time, relative to one another. Based on recent evidence pointing to domain-general reward processing in social and non-social decision-making (Arabadzhiyska et al., 2021; Rilling and Sanfey, 2011; Wake and Izuma, 2017) and previous reports showing fMRI correlates of risk in social contexts (Vavra et al., 2018; Xiang et al., 2013), we hypothesized that socially-relevant risk and RiPE signals would manifest similarly to previous non-social tasks.

More importantly, if signed RiPE signals are indeed operationalized as the difference between trial-specific risk estimates and the overall expected risk, then it can be predicted that momentary signals of experienced risk would be computed and emerge first, followed by a later signature of signed RiPE. While signed RiPE signals are thought to be involved in updating risk estimates (Preuschoff et al., 2008; d’Acromont et al., 2009), unsigned RiPE likely control the speed of risk updating similar to the hypothesized role of unsigned reward PEs in controlling the learning rate (the weight applied to

the signed reward PE) in a reinforcement-guided framework (Mackintosh, 1975; Pearce and Hall, 1980). As such, unsigned RiPEs could emerge near simultaneously with the signed RiPE signal to drive adaptation to risk and learning, as shown previously for reward updating (Fouragnan et al., 2015, 2018, 2017).

2.2 Materials and Methods

2.2.1 Participants

The final data reported here includes 40 participants (mean age: 25; SD = 9.19, female/male = 27/13). The participants had no diagnosed neurological or psychiatric conditions and were asked to not consume alcohol or other psychoactive substances the day before and on the day of the experiment. EEG data were originally collected for 71 participants recruited through the online system for experiments in psychology at The University of Glasgow. Eight of them were not used for the analysis as they only took part in a pilot task that was significantly changed for the main experiment. Five of them were excluded because of failure to report back the outcome of a trial in more than 10% of trials. Four of them were excluded due to the excessive eye and/or movement artifacts in the EEG data. The remaining 54 participants were included in our multivariate single-trial linear discriminant analysis (LDA, see Methods section Single-trial EEG analysis). Due to the nature of the task, which does not require an active decision, or any behavioural response relevant to the selection of the stimuli (see the Main task), we expected that some participants may fail to pay attention and/or encode the relevant stimulus features, which would ultimately manifest as poor discrimination performance along the relevant stimulus dimensions (e.g. low-vs-high risk/RiPE/|RiPE|). We accounted for this initially by recruiting more participants than would normally be sufficient for similar EEG studies and analysis techniques (Fouragnan et al., 2015, 2017; Philiastides et al., 2010b). Furthermore, EEG discrimination performance for 14 participants was not statistically different from chance, for

all three signals of interest (as explained under the ‘Bootstrapping-based exclusion’ section), and they were therefore removed from the final analysis/results. The study was approved by the College of Science and Engineering Ethics Committee at the University of Glasgow (300210143) and informed consent was obtained from participants prior to the experiment.

2.2.2 Risk and RiPE estimation

During the main experiment, participants were presented with pairs of faces that were associated with different social value estimates (£50/70, £40/80, £20/100, £60/60), one of which was selected as the outcome of each trial. Figure 2.1 shows the spread of values around the EV and the size of current risk and RiPE for each pair. Table 2.1 and Figure 2.3 show the values that resulted from the calculation of each risk variable of interest following equations 2.1 to 2.3. The equations provided show the way we calculated the expected value (EV – Eq. 2.1), the SD of all rewards (predicted risk – Eq. 2.2), the difference between each reward and the average reward (current risk), and the deviation of each reward from the SD (risk prediction error – Eq. 2.3), following the approach of O’Neill and Schultz (2013). For each participant, we computed the current risk, and the signed and unsigned RiPEs, which were used in our single-trial participant-specific LDA analysis.

$$EV = \sum_{i=1}^n (p_i \cdot x_i) \quad (2.1)$$

$$SD = \sqrt{\sum_{i=1}^n p_i \cdot (x_i - EV)^2} \quad (2.2)$$

$$RiPE = \text{Current } SD_{Risk \text{ cue}} - \text{Predicted } SD_{Fixation \text{ spot}} \quad (2.3)$$

Where n is the number of rewards, p is the probability of each reward, and x is the amount of each reward in points. The Expected value (EV, Eq. 2.1) is the average of all rewards. The current risk is the standard deviation of a pair of rewards on each trial. For example, for the 50/70 pair the current risk would be $(70 - 50) \cdot 0.5 = 10$. The ‘risky’ pairs are 50/70 (low risk), 40/80 (medium risk), 20/100 (high risk). The predicted risk is the SD of all rewards in the experiment (Eq. 2.2). As can be inferred from equation 2.3, due to the fixed value of the predicted SD across the whole task, it follows that the signed RiPE is correlated with the current risk (trial risk), and the two signals cannot be distinguished. To tackle this and obtain a measure of RiPE that is decorrelated to current risk, we introduced a set of trials in which the two possible rewards are the same and equal to the average of all rewards in the task (60/60). Hence, there was no variability, and the current risk was 0. This resulted in different values for each risk variable on each trial set, as shown in Figure 2.2, Figure 2.3, and Table 2.1. Moreover, we introduced a different colour fixation cross to indicate a 60/60 trial was entered, leading to the predicted SD at the fixations cross on these trials being different from the overall predicted SD.

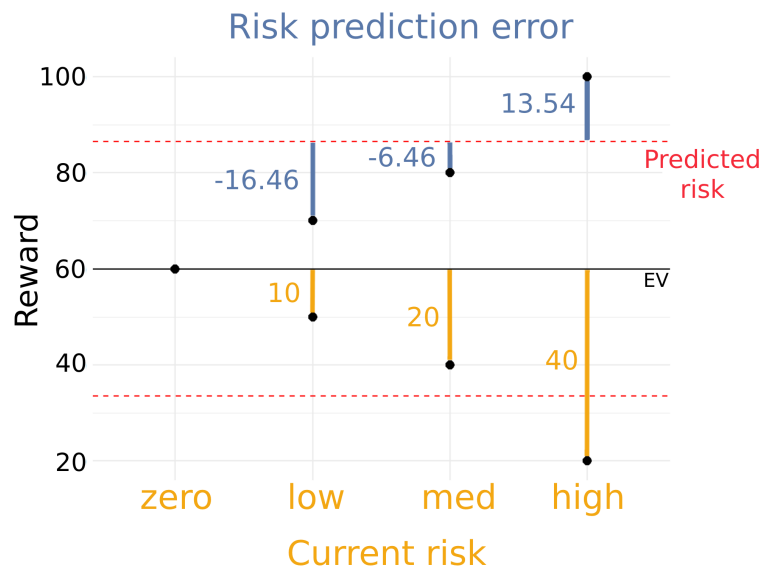


Figure 2.1 Values of pairs of stimuli. Current risk is shown in yellow, RiPE is shown in blue, and predicted is, which is constant within risky trials is shown with red.

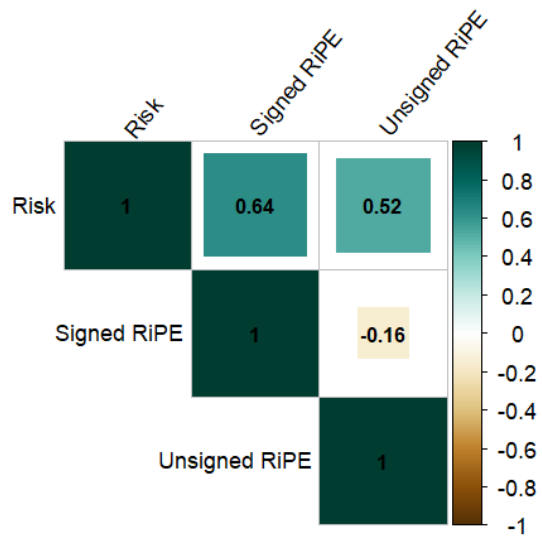


Figure 2.2 Correlation plot of computed values.

Pair category	60/60	50/70	40/80	20/100
Predicted risk	26.46	26.46	26.46	26.46
Current risk	Zero risk = 0	Low risk = 10	Medium risk = 20	High risk = 40
Signed RiPE	0	-16.46	-6.46	13.54
Abs. RiPE	0	16.46	6.46	13.54

Table 2.1 Computed values of each risk condition across trials. Shading represents LDA training and testing levels. Blue shading shows each condition's lowest (light blue) and highest (darker blue) values.

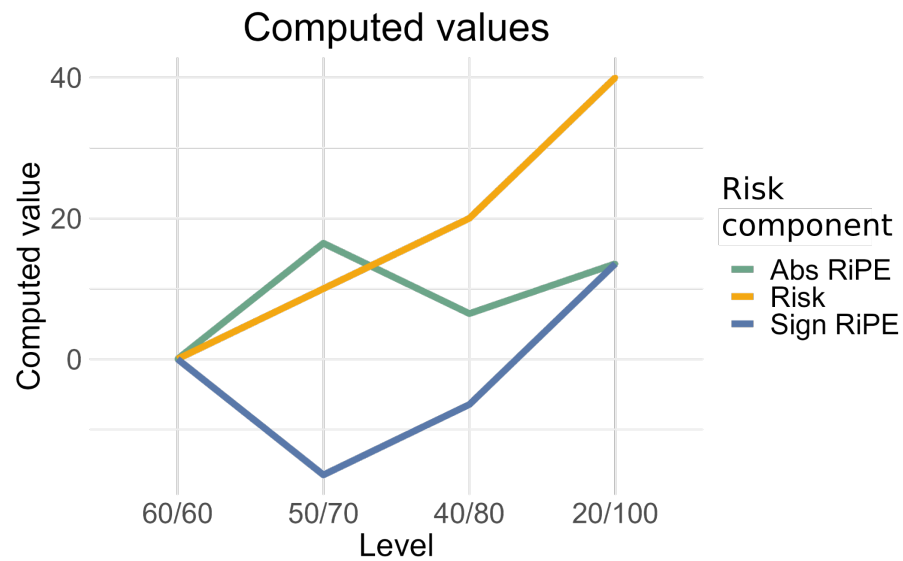


Figure 2.3 Computed levels of risky and safe trials. The figure shows that the three components can be distinguished.

2.2.3 Stimuli and task

The task was performed in an electrically shielded, soundproof booth with a small amount of light from the ceiling. The stimuli were presented on an Nvidia 4k monitor with a 120Hz frame rate, approximately 60 cm away from the participants. Psychopy software v2020.2.3 (Peirce et al., 2019) was used to present the task and record behavioural responses. The task was presented on a black background (Figure 2.5), and each pair of faces was centred in the middle of the screen. Responses were made with a small gaming controller with 5 buttons connected to the computer with a USB cable.

Unlike previous work using abstract stimuli (i.e., non-social) to create associations with expected rewards expressed as money or points, here we used socially relevant stimuli obtained in the context of an economic game (i.e., trust game). Specifically, one day before the experiment, each participant rated 28 photo-realistic images of faces (see Figure 2.4. Example stimuli, Arabadzhiyska et al., 2021; Yu et al., 2012; Zhan et al., 2019) in 5 blocks. The rating was based on a hypothetical scenario in which participants were told the faces were photos of people who played in a Trust Game of two players

(an investor and a trustee). Participants were told that all trustees received £200 from the investor and returned between £10 and £120, and the participant's task was to estimate how much each trustee returned. Hence, we consider the participant's subjective judgments of trustworthiness as a proxy of the economic outcome value in the context of a one-shot money-incentivised game (Eckel and Wilson, 2003; Scharlemann et al., 2001). To measure consistency, we used the standard deviation across the ratings of individual faces across 5 blocks. From the 28 rated faces, for each participant, we selected those with the most consistent ratings (lowest SD) on each of the points ranging from £10 to £120, in increments of 10, resulting in 12 faces. Further, instead of using all 12 faces per participant, we chose 7 faces which would represent values that allow decorrelation between the variables of interest, hence, we chose one face associated with each of the following values: £20, £40, £50, £60, £70, £80, £100. With this method, we fulfilled our aim to obtain 7 faces for each participant which ranged from least to most trustworthy, being associated with low or high amounts of points respectively, and spanning the predetermined scale of rewards. For the main task, we represented the pounds as points (50/70, 40/80, 20/100, 60) to be able to motivate the behaviour to perform correctly and win 'points' that were later associated with a final payout (see Main task).

All faces were adjusted for height and width in such a way that when presented in a pair with any of the other 6 faces, there was no striking contrast between the two faces. All faces ranged from 212 to 240 pixels in height and 136 to 180 pixels in width and were presented on a black background (stimuli were up to degrees of visual angle: 2.82° height and 2.11° width).




Figure 2.4 Example stimuli. Faces of various trustworthiness levels were used in the task. Each participant initially rated 28 images of faces framed in the context of an economic game. (See 2.2.3).

2.2.4 Main task

With a task preceding the main experiment, we aimed to strengthen the association between the ratings each participant gave to each of the final 7 faces. As the faces were rated at home one day before the experiment and were the final selection from a sample of 28 faces, it was important to make sure the participants remembered well the ratings they gave specifically for these 7 faces. The task consisted of each participant observing their preselected set of 7 faces one at a time, presented repeatedly in a randomised order while trying to remember and indicate the value of each face in points using a scale with the 7 possible rewards (20, 40, 50, 60, 70, 80, 100). Each face was presented in the middle of the screen with a discrete scale of 7 points. After selecting a point on the scale, the participants received the message ‘correct’ or ‘wrong’ coloured in green or red, respectively. As such, the task prompted learning through trial and error and was finished after 14 consecutive correct answers. All participants succeeded in memorising the faces in less than 10 minutes.

After memorising the values associated with the 7 faces, the main task was explained to the participants, and more specifically the way the faces would be paired in relation to their risk levels. During each trial, one face from a pair was highlighted and the points won on that trial were the points associated with the selected face (as learned during the preceding consolidation task). It was further explained to the participants that there were overall seven faces, which were divided into pairs in such a way that

each pair formed an opposition along the middle of 60 points. The faces representing 50 and 70 points were paired together, forming a pair with ‘low’ risk, the pair of faces giving 20 or 100 points was categorised as ‘high risk’, and the ‘medium risk’ pair was that of faces of 40 and 80 points. There was also a ‘zero’ risk condition in which a pair was formed by the same face shown twice (60/60). After the participants understood the possible pairs, they were shown an example of the main task.

The main task is represented in Figure 2.5. Each trial began with one of two cues indicating if the current trial would be of ‘zero’ risk with the same faces, or a trial with any of the other risky pairs. Risky pairs (indicated by a cue with a blue line around all risky values) were any pairs that consisted of two different faces, whereas safe pairs (indicated by a cue with a pink line around the average value) were pairs consisting of the same face shown twice. We decided to use a different cue for zero-risk trials to completely remove the risk prediction error. As the cue prepared the participant for a trial with no risk, the overall standard deviation was expected to be reset to zero, leading to no difference between expected and experienced risk when the facial cues were presented. Otherwise, as in the risky set of trials, it can be expected that there is a risk prediction error computed as the difference between the overall standard deviation and the current risk of 0.

The cues stayed on the screen for a variable amount of time drawn randomly from within the range of 1 second to 1.4 seconds. A pair of faces appeared on the screen for a time in the range of 1.5 to 2 seconds, which is the period we expected risk and RiPEs would be encoded (Lauffs et al., 2020; O’Neill and Schultz, 2015). After that one of the faces was highlighted by an orange frame indicating the reward outcome (as inferred by the selected face). The orange highlight frame stayed on-screen for a variable duration between 0.8 and 1.25 seconds. Following this was a screen with a graphical representation of two scales that aimed to prompt participants to pay attention to which face was selected. Participants had to remember the value that was won from the face selected on the previous screen and indicate both the level of risk that the face

was part of (zero, low, medium, or high) and the specific value (20, 100, 40, 80, 50, 70, or 60). The risk was selected first by moving a red dot left or right until the desired level was indicated by the dot. Then a side button on the controller was pressed to move to the second scale, where the dot automatically moved to the long horizontal line representing the average and right in the middle of the two values associated with the previous selection. So, for example, if the low risk was selected the red dot was then moved to the middle between values 50 and 70. Next, the participants had to press the ‘up’ or ‘down’ buttons to make their final selection about the value they believed they won. For example, they would press ‘down’ to select 50, and ‘up’ to select 70. In the case of ‘zero’ risk, there was no active selection on the scale apart from participants pressing the same confirmation button twice: once for the first scale, selecting ‘zero’ risk, and a second time to confirm the value of 60. The red dot was first shown on the ‘risk’ scale and then in the middle of the ‘value’ scale. This was always the case, independent of whether the left or right face was highlighted previously as the face was the same and carrying the same reward of 60. With this procedure, there was no difference in the number of buttons being pressed during the risky or zero-risk conditions. After the value selection (or confirmation in the case of zero risk) was made, the trial finished, and a new cue appeared on the screen.

There were 60 trials per risk level. This resulted in 240 trials per person, separated into 15 blocks with breaks in between. On each trial, if there was no response on the first scale by the participant moving the red dot left or right in less than 2 minutes, a message ‘Too slow!’ appeared on the screen and the next trial began. The whole task (as shown in Figure 2.5) took approximately 35 minutes per person.

The incentive for participants to strive to be correct (and pay attention) on all trials was their knowledge that after the experiment had ended, a certain number of trials were chosen randomly, and if their response was correct on all the chosen trials, the final payment would be increased by the number of points indicated on those trials, resulting in payment between £12 and £18. All participants were debriefed and received payment according to the time spent in the lab (approximately £16).

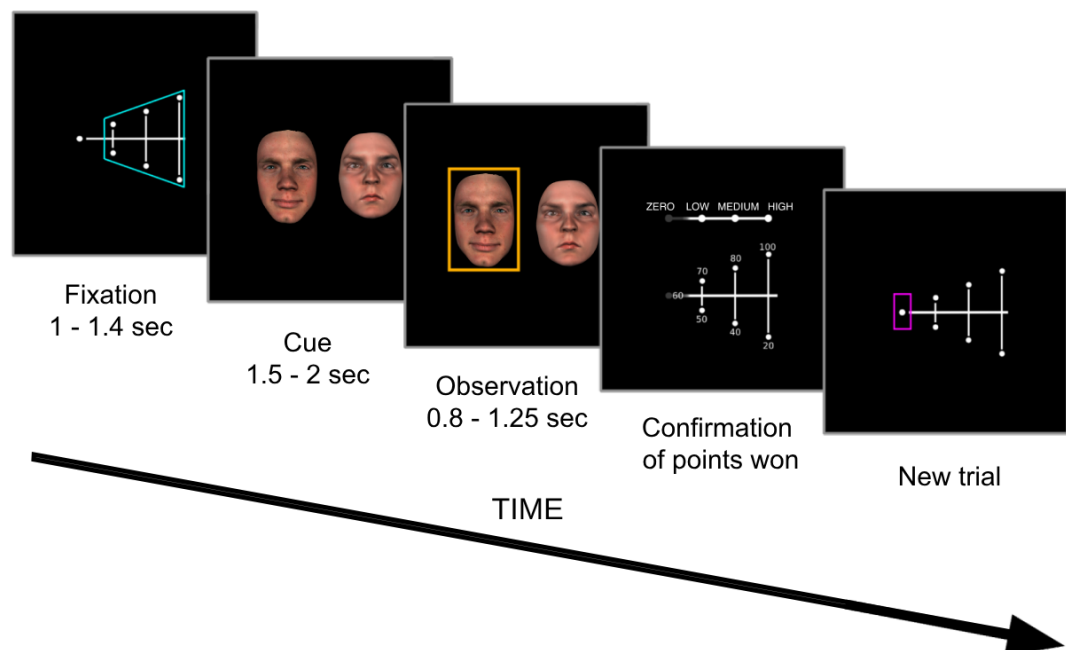


Figure 2.5 Experimental timeline. Each trial began with one of two cues indicating if the current trial would be of ‘zero’ risk with the same faces, or a trial with any of the other risky pairs. The cues stayed on the screen between 1 and 1.4 seconds. A pair of faces appeared on the screen for a time in the range of 1.5 to 2 seconds, after which one of the faces was highlighted by an orange frame lasting for 0.8 to 1.25 seconds. A graphical representation of all values appeared on the screen, which prompted the participants to select the value they won (known from memory, see 2.2.4). After selection, a new trial started.

2.2.5 EEG data acquisition

An EEG system from Brain Products was used with a cap with 64 Ag/AgCl scalp electrodes placed according to the international 10–20 system on an EasyCap (Brain Products GmbH, Germany, EEGLAB BESA) with a ground electrode on the chin, reference electrode on the left mastoid, and NFCz electrode in the middle of the forehead. The impedance of each channel was lowered below 20 kOhm, if possible, and to a maximum of 50 kOhm. The EEG recording computer received experimental event markers from the machine presenting the stimuli with Psychopy, via a parallel port. Sampling was done at 1000 Hz with BrainAmp DC amplifiers (Brain Products, Gilching, Germany).

2.2.6 EEG data pre-processing

During recording, data were filtered online with an analogue band-pass filter of 0.0016 Hz to 250 Hz. Further offline pre-processing of the individual EEG data was done in Matlab (ver. 2019b) and plotting and further analysis of the LDA results were done in Matlab and R (R Team, 2021). The raw EEG data were explored via EEGLAB v.2021.1 (Delorme and Makeig, 2004) in Matlab, identifying ‘bad channels’. Individual channels were removed if they were with consistently high noise, observed as fluctuations throughout the whole experiment period, and selected visually with EEGLAB. No more than 5 channels were removed per participant. Moreover, we removed trials in which the amplitude of the average signal across all channels was higher than the absolute of 2.56 SDs above the mean. The maximum number of trials removed was 5. For filtering the data further and removal of eye movement artifacts, we used an in-house pre-processing pipeline (Franzen et al., 2020; Gherman and Philiastides, 2014). Filtering the raw data included band-pass filtering at 0.5 - 40 Hz and re-referencing data to the average. The presentation of the face pairs was the stimulus onset of interest (second screen in Figure 2.5), the timing of which was informed by the time-precise

event markers recorded in the raw EEG data. The filtered EEG data were epoched by taking 100 ms prior to and 800 ms after the onset of the face pairs for all trials and electrodes. The data were baseline corrected, defining the baseline interval as 100 ms before stimulus onset. Each participant completed an eye-calibration task which included blinking repeated for a few seconds while fixated on a cross in the middle of the screen and following a white cross on a black background moving up-down, left, and right to the edges of the monitor. Further, linear components (sensor weights) associated with eye blinks and horizontal and vertical eye movements were identified using PCA and projected out of the EEG data from the main task (Franzen et al., 2020; Parra et al., 2003).

2.2.7 Single-trial EEG analysis

We analysed the EEG data with single-trial multivariate linear discriminant analysis (LDA) to obtain components that capture the neuronal activity with a better signal-to-noise ratio compared to the original activity from single electrodes (Franzen et al., 2020; Parra et al., 2003; Philiastides and Sajda, 2006; Sajda et al., 2009).

First, for each participant, all trials were divided into 4 levels or classes (as shown in Table 2.1), according to the four possible values that each condition could take. We selected only half of all trials ($n = 120$). In this way, we trained the classifier on part of the data, namely, the two classes: highest and lowest per condition (Table 2.1 – dark blue shading). We later tested the results on the rest of the trials (Table 2.1 – light blue shading; explained under Identifying risk-related EEG components).

Finally, we performed three separate LDAs, aiming to examine the temporal and spatial profile of the three conditions. The EEG data is denoted as $x(t)$, that is the multidimensional data, with D channels, transformed into a single vector at time t . It is indexed by the two classes of trials and weighted by a spatial filter (w) to produce a one-dimensional discriminating component $y(t)$:

$$y(t) = w^T x(t) = \sum_{i=1}^D w^i x^i(t) \quad (2.4)$$

T refers to the transpose operator and D refers to the number of EEG sensors. The resulting spatial weighting vector $y(t)$ is a linear projection of the data, serving as a decision boundary, maximally separating the EEG data according to the two pre-defined classes, while minimizing the variance within the data corresponding to each class. The discriminator maps amplitudes $y_i(t)$ for the lowest and highest levels of each component onto negative and positive values, respectively. The optimal w is obtained by a discriminant function, formulated with logistic regression, determining the log-likelihood of each class given the data $x(t)$, and an optimization algorithm based on iteratively re-weighted least-squares (IRLS, Parra et al., 2005; Philiastides et al., 2010b). Further, the spatial filter (w) was obtained by a sliding window approach in which each trial of the epoched data is separated into short windows of 60 ms, centred around every 10 ms from 100 ms before to 800 ms after the stimulus onset (-100 to 800). To quantify the LDA performance, a Receiver Operator Characteristic (ROC) analysis was done, with an area under the curve (AUC) denoted by A_z . A Leave-one-out cross-validation was used to avoid overfitting by using N-1 trials for training the spatial filter and testing it on the left-out trials for all sliding time windows. A_z values range from 0, through 0.5 (chance performance), to 1 (perfect separation), and indicate the probability of trials being correctly mapped to the extreme levels of each component. To approximate the sources of the discrimination for each component, we used a forward model found by linearly predicting $x(t)$ from $y(t)$, resulting in the spatial distribution of the normalized correlation between the data and the component:

$$a = \frac{X \cdot Y}{y^T y} \quad (2.5)$$

This resulted in a which is a scalp projection subsequently plotted as a scalp topography or a heatmap.

2.2.8 Bootstrapping-based exclusion

We performed bootstrapping for each participant's Az values to obtain an Az level that would lead to an alpha cut-off of 0.05. The leave-one-out test was performed after randomizing the truth labels of the discriminating classes 1000 times (previously shown as a sufficient sample (Rousselet et al., 2019), resulting in a probability distribution with Az values. For each participant, we found a bootstrapping threshold, which was the cut-off Az value at 5% of both tails of the distribution, and then selected all Az points that were above this value. For some participants, on all three components, there were no significant Az values that were above their threshold in the time window 0 ms to 700 ms locked to stimulus onset. We considered this as a sign that the participant failed to pay attention to the task, failed to encode any of the variables, and were excluded from the analysis.

2.2.9 Identifying risk-related EEG components

The LDA analysis was run on each risk variable (current risk, absolute RiPE, and signed RiPE) separately. As shown in Table 2.1, on each trial, the risk variables can hold one of four values, associated with the rewards on the current trial. To ensure the LDA-resulting weights were not overfitting to the extreme conditions, we introduced cross-validation by dividing trials within each risk component into training (Table 2.1 – dark blue shading) and testing data sets (Table 2.1 – light blue shading). We applied the spatial filter (w) on the 'unseen/testing' data consisting of the middle values for each risk component (light blue shading on Table 2.1). For current risk, the lowest value was 0 (zero) while the highest was 40 (high risk), the middle two values 10 (low risk) and 20 (medium risk) were not used while training the classifier. For signed RiPE, the lowest value was -16.46 (high negative), and the highest was 13.54 (positive), while the middle values were -6.46 (low negative) and 0 (zero RiPE). The classifier for absolute RiPE was trained on 0 (zero abs. RiPE) and the highest value of 16.46 (high abs. RiPE), while the

middle two values comprised 6.46 (low abs. RiPE) and 13.54 (medium abs. RiPE). We hypothesised that if the spatial filter of the data, y , is indeed specific to a component, then we would observe a parametric effect across all 4 levels of the component, from lowest to highest. This effect is quantified by a significant difference (in the predicted direction) between the y 's at the middle two values. We used a paired-samples t-tests for each risk condition, reporting p-values, mean differences and confidence intervals (as difference = x 95% CI [low CI, high CI]). We further supply the results of a Bayes Factor test with paired samples, obtained using the R function 'ttestBF'.

After obtaining the group-averaged Az values, we considered a broad window from 100 ms to 700 ms post-stimulus-onset, in which Az values were above chance performance (0.5). We average the forward model estimates (a , Eq. 2.5) across participants at each timepoint (every 10 ms) across the 100 ms – 700 ms offset window. Because for each participant a risk component can occur at a different timepoint, we aimed to find this timepoint per participant that is most consistent with the component of interest. We could have picked the highest Az point per participant; however, the highest discrimination performance may not be indicative of when the component occurs. Instead, we looked for time points at which, first, the difference between the y 's at the two middle values was highest and, second, a timepoint at which the scalp projection was closest to the source representation (the group average a).

First, for each participant, we computed the difference between the two y -values of the 'unseen' data and selected the top 20% of time points with the largest differences, independent of the direction of the difference (i.e., regardless of whether it was in the predicted direction). Importantly, the time points in which we looked for the largest differences were within the 100 ms to 700 ms window. Second, to obtain an average scalp topography (map) for all participants, instead of averaging across an arbitrary window around the peaks, we performed a k-means clustering analysis to define the cut-off points of distinct scalp topographies around the average peaks. We used a k-means clustering algorithm with a squared Euclidean distance metric on the

intensities of a vector of group-averaged scalp topographies (a's) across offset times. The time range was different for each risk component and was limited by 100 ms prior to the first prominent peak and 100 ms post the final prominent peak appearing before 700 ms post-stimulus onset. The number of different time windows with similar scalp topographies, k , was optimised using silhouette values via the MATLAB `evalcluster` function.

For example, as illustrated in Figure 2.6, we performed k-means on a window between 200 – 500 ms post-stimulus onset, because this window encapsulated the group peaks. For signed RiPE the most consistent topographies around the highest Az values started from about 280 ms and continues to 460 ms post-stimulus onset (yellow and green clusters, Figure 2.6). In this example, we would average the scalp maps within the window 270-460 ms and use this group average for further comparisons with individual-specific scalp maps.

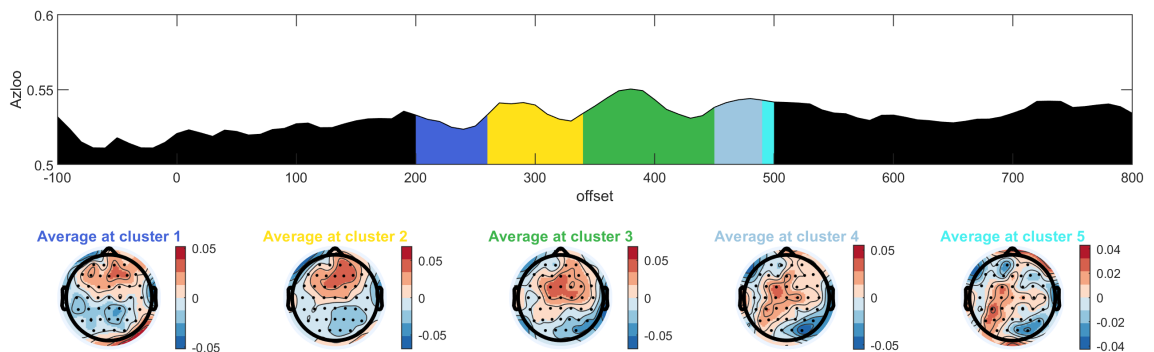


Figure 2.6 Example of results of k-means clustering of average scalp topographies. Each topographic map is the average of all maps within the same-coloured window on the Az plot.

From the top 20% largest differences in y 's at the middle two levels of each risk condition (as outlined above), we chose a time point with a scalp topography that had the best correlation with the cluster-derived average of all participants.

The above procedures resulted in a single point per participant that was considered the ‘best’ time point to represent the component. We ran a t-test on the middle levels (testing set trials, light blue shading in Table 2.1) using the y-values at the best time points to test the hypothesis that there would be a significant difference in y-amplitudes between the middle levels, showing the expected directional parametric modulation of each risk component (as in Figure 2.3). All the above procedures were done separately for current risk, signed RiPE, and absolute RiPE. Further, we compared the average time points per risk condition via a density plot of stimulus-locked offset times.

2.3 Results

The main aims of this work were to test for evidence for the three risk-related signals that were proposed to be encoded by the brain, whether this can further be achieved given our socially relevant stimuli, and if these components were present, to determine the spatio-temporal profile of each one.

The cross-validation showed that each component was capturing EEG activity specific for the related risk variable, as predicted by the relationship between raw computed values as shown in Figure 2.7. Figure 2.8 shows the discrimination performance for each risk component which for all three signals was above chance level (0.5 Az leave-one-out).

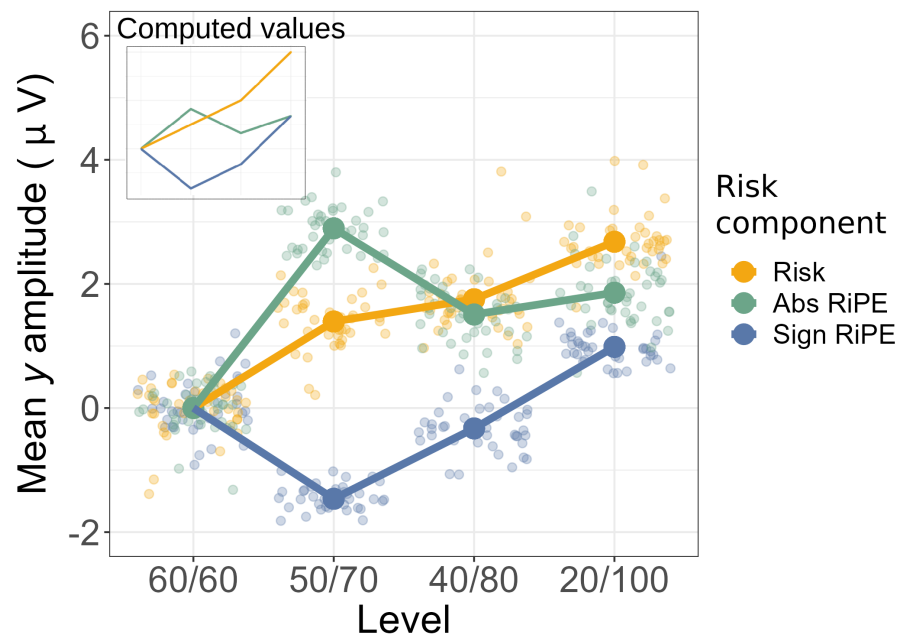


Figure 2.7 EEG-derived values of each risk variable across the four levels of trials. Top left window shows the original computed values (same as Figure 2.3). Dots are individual participants.

2.3.1 Current risk

The discrimination between the zero and highest risk had an average peak at 272 ms (max AUC = 0.59) and continued to decline gradually, as shown on the top row of Figure 2.8. The best time points per participant were obtained as outlined in the Methods section, and we did an analysis of the difference between the values of the discriminating components, y , for the middle two conditions of current risk. As shown in Figure 2.9 (left), the y values at each level of current risk are consistent with the expected parametric effect as shown in Figure 2.3. The difference between the y -amplitudes at the two testing sets of trials was significant (paired $t(39) = 4.297$, $p < 0.001$, difference = -0.36, 95% CI = [-0.53, -0.18], BF = 220). Hence, the component amplitude at the medium current risk condition was significantly different from the low current risk (in 75% of participants), confirming the parametric relationship between the levels of risk and y -amplitudes. Shown in Figure 2.8 (top row, right), the average scalp projections from the discriminating component for current risk indicate a significant difference between lowest and highest current risk values with a positive correlation at left fronto-central electrodes and right frontal electrodes (positive correlation) and a negative correlation at parietal sites of the right hemisphere.

2.3.2 Absolute RiPE

The unsigned/absolute RiPE discrimination (middle of Figure 2.8) on average peaked at 280 ms (max AUC = 0.59). There was a significant difference in discriminator amplitudes between low and medium unsigned RiPE ($t(39) = 3.9$, $p < 0.001$, diff = -0.35 [-0.53 -0.17], BF = 75.9). As shown in Figure 2.9 (middle), the average y -amplitude at zero absolute RiPE ('60/60') was higher than for the lower value of absolute RiPE that was at trials '40/80' (for 80% of participants).

The average scalp projection from participant-specific best peaks shown in Figure 2.8 middle row, is expressed in a negative correlation at fronto-central electrodes, and a positive correlation at parietal electrodes.

2.3.3 Signed RiPE

The discrimination of the lowest and highest values of signed RiPE (bottom of Figure 2.8) contained multiple peaks between 270 and 480 ms, and 730 ms post-stimulus, indicating greater temporal inter-participant variability (max AUC = 0.55). For this component, we also found a significant difference after projecting the classification vector, w , onto the middle two values for signed RiPE. Specifically, the average y amplitudes at low negative signed RiPE trials ('40/80') was lower than the y amplitudes at trials with zero signed RiPE ('60/60') in 75% of participants, $t(39) = 5.9$, $p < 0.001$, $\text{diff} = -0.33 [-0.44 -0.22]$, $\text{BF} = 25412$, also shown in Figure 2.9 - right.

The average scalp topography from best points at the level of signed RiPE (Figure 2.8 bottom right) showed an opposite pattern to unsigned RiPE, with a similar map but a flipped dipole. The positive correlation is observed at fronto-central electrodes, while a negative correlation occurs at occipito-parietal sites. This tendency suggests that the source of unsigned and signed RiPE may be similar and with the reversed direction of activation.

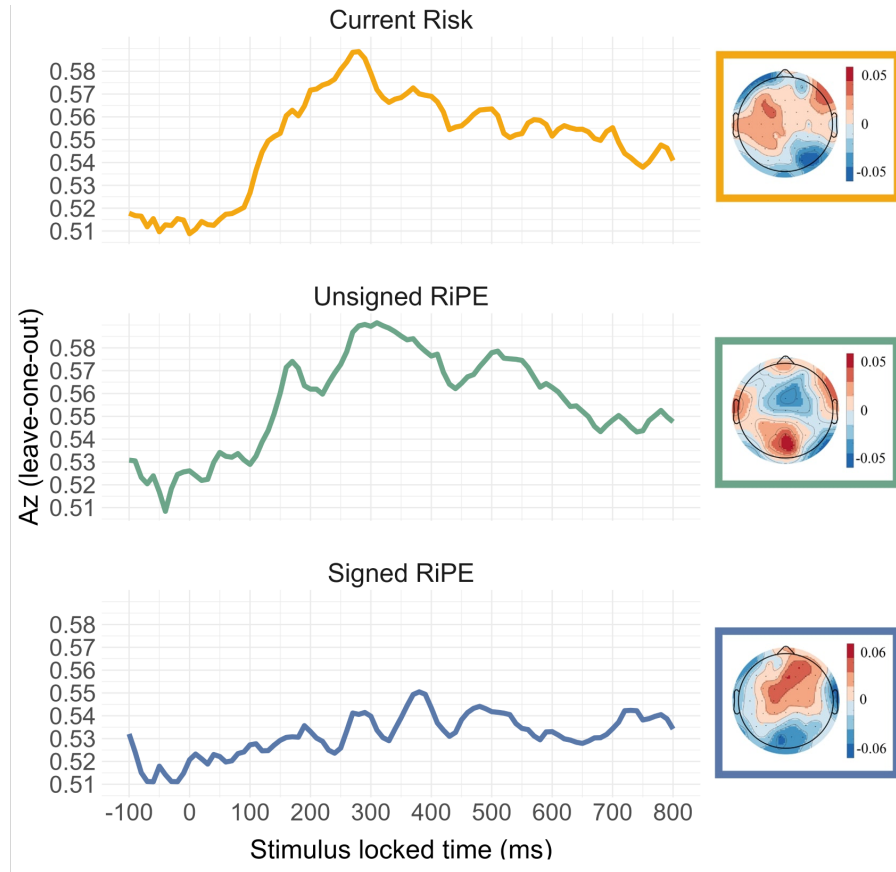


Figure 2.8 Average discrimination ($n = 40$) and average topographies from best discrimination points.




Figure 2.9 EEG amplitudes at all three conditions derived from trials used for training the LDA (pink) and testing the resulting weights (black).

We also plotted the onset times of the best points derived from individual participants as density plots as shown in Figure 2.10. Although there are big overlaps and no significant difference between the three components, generally current risk is computed earlier, followed by unsigned RiPE, and then by signed RiPE.

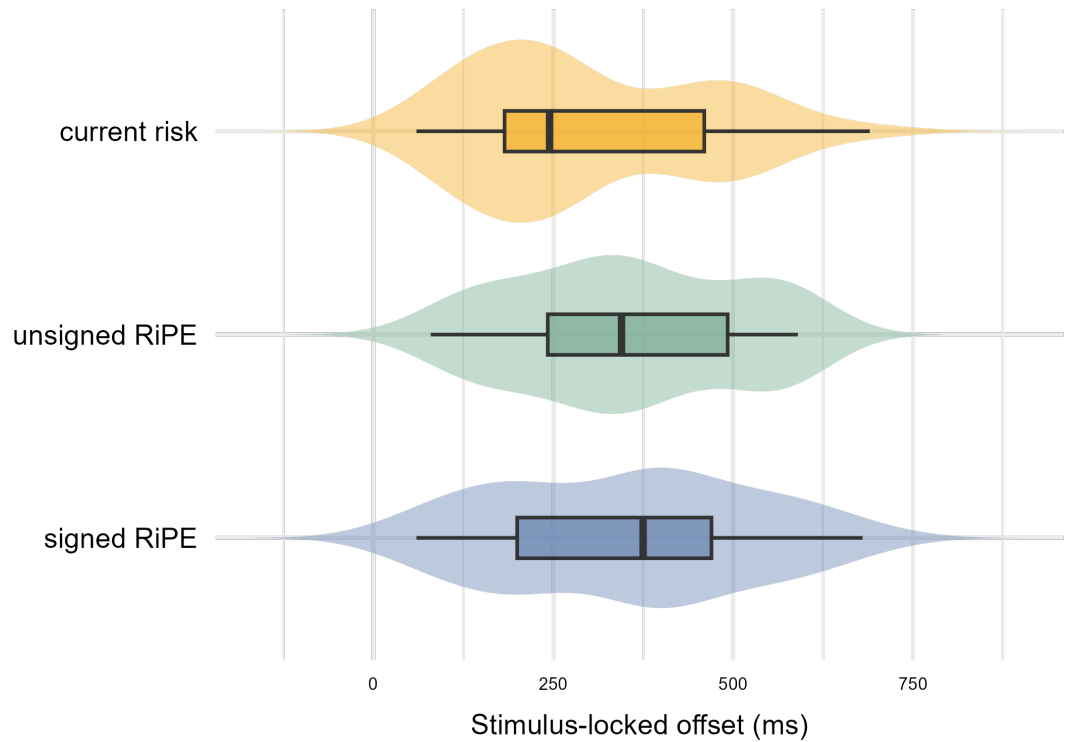


Figure 2.10 Density plots of the time of onset of best points relative to the onset of the stimulus.

2.4 Discussion

We used a passive-observation task with socially relevant individual-specific cues as rewards and proposed that risk coding and risk prediction errors will be reflected in the EEG signal like previous non-social probabilistic tasks. Namely, we hypothesised that if the risk is encoded separately from rewards and follows a reinforcement learning profile like reward coding, we would find distinguishable components in terms of time and EEG space. Our results show that predicted risk, signed and unsigned RiPEs were indeed separable with EEG. First, we found spatio-temporal differentiated components associated with risk coding, unsigned RiPE, and signed RiPE. Second, we showed current risk coding to appear earlier than unsigned RiPE, followed by signed RiPE. Finally, the above EEG components results were validated by establishing their monotonic relationship with the levels of the three risk-related variables. These findings are significant for the current efforts to delineate risk processes and relate them to learning in complex uncertain environments.

In previous fMRI studies using financial uncertainty, perceptual uncertainty, and the balloon analogue risk-taking task, risk correlates were found in the anterior insula, the OFC, inferior frontal gyrus (IFG), the lingual gyrus, or the posterior cingulate cortex (Burke and Tobler, 2011; Huettel et al., 2006; Payzan-LeNestour et al., 2013; Preuschoff et al., 2008; Schonberg et al., 2012; Tobler et al., 2007). The above results are consistent with lesion studies showing increased gambling, or maladaptive risky choices in patients with insular or OFC lesions, while reward estimates were not impaired (Clark et al., 2008). Our results confirmed these previous findings by showing that by keeping average rewards constant, a brain response can be delineated that varies with different levels of variance of the reward distributions. Therefore, in terms of the separation between predicted rewards and risks, our results suggest that risk is separately coded from rewards and that this can be tracked by EEG.

Two fMRI studies found RiPE correlates in the anterior insula, inferior frontal gyrus, and ventral striatum (d’Acremont et al., 2009; Preuschoff et al., 2008). In the duration of 8 seconds post-stimulus, using fMRI, Preuschoff et al. (2008) showed that RiPE signals appeared fast after stimulus onset (2 sec), whereas, predicted risk signals were delayed (8 sec), which is contrary to our findings showing risk coding to appear first followed by RiPEs. These discrepancies may be due to the different imaging modalities being implemented in the two studies or due to the differences in experimental design, especially given the long delays between trials in the fMRI study. Previous animal studies showed that risk and RiPE signals take less than 1 second post-stimulus and have an overlapping time signature (Fiorillo et al., 2005; Kobayashi et al., 2010). Hence, fMRI, which depends on blood oxygenation levels and has a poor temporal resolution, is ineffective in detecting such fast brain responses. Our analysis of EEG data showed that risk components can be detected as fast as 800 milliseconds post-stimulus onset, with peak onsets close in time but distinct. Using a brain imaging device with high temporal resolution, such as EEG, is essential in distinguishing between these signals that appear to arise almost simultaneously within the span of a decision-making trial.

An EEG study by Lauffs et al. (2020) used card stimuli to cue perceptual risk as participants did not receive a reward depending on the values of the cards but rather on their success in predicting if the second of two cards presented in succession would be higher or lower. The cards ranged from 1 to 10 in value, the risk was measured as the variance of the reward distribution, while RiPE was defined as the squared deviation of the reward prediction error from the variance. This type of design aims to measure the perception of the numerosity of sensory stimuli (Kobayashi et al., 2010, allowing depiction of brain processes for internally represented objective mathematical variables. The results showed that there were two components following the presentation of each card, the earlier representing the salience of the stimulus, while the later one representing the RiPE. Our results are consistent with the time of build-up of the RiPE signal, which in the study by Lauffs et al. (2020) started around 300 ms post-stimulus onset, and for the current study times of highest discrimination were 300 ms and 380 ms post-stimulus

for unsigned and signed RiPE, respectively. Moreover, after source reconstruction, the authors found the saliency-related ERP to be associated with visual cortices, but also with the insula which was the main source for the second ERP component, related to RiPE. This is consistent with the spatial projections of the components we found for unsigned and signed RiPE, with the two showing flipped dipoles (in terms of sign), but similar scalp topographies. This finding suggests that the two components may originate in the same source but have a different direction of propagation of the signal.

The parametric effect we found in the current risk component is consistent with previous studies with monkeys showing increased firing of single neurones with increased levels of uncertainty (Grabenhorst et al., 2019; O’Neill and Schultz, 2015), and the monotonic effect of the components associated with unsigned RiPE is similar to previous animals studies (Kobayashi et al., 2010; O’Neill and Schultz, 2013) as well as fMRI and EEG findings (Lauffs et al., 2020; Preuschoff et al., 2008). A previous animal electrophysiology study, on which we based our paradigm of risk, showed risk and unsigned RiPE coding in OFC single neurones of rhesus macaques (O’Neill and Schultz, 2013). In the study, however, the signed RiPE was correlated with ‘current risk’ (trial risk) that was subtracted from predicted risk (overall task risk). As predicted risk was constant across the experiment, this meant that current risk and signed RiPE could not be separated. This makes our study the first one to show a parametric effect of the signed RiPE that is dissociated from both risk and unsigned RiPE.

We showed that the component associated with risk appeared before both signed and unsigned RiPE components, which is consistent with the suggestion that risk processes are separate from reward evaluation but may follow a similar pattern. The temporal profile of the three signals is expected. As the RiPE is derived from the evaluation of current risk, it can be assumed that risk is required to occur first in order to update the prediction error values. Similar to rewards, the unsigned RiPE is likely equivalent to a surprise signal, influencing the adjustment of a learning rate. While the signed RiPE aids the adjustment of a value function signifying whether predictions are better or

worse than expected (O’Neill and Schultz, 2010a). The unpredictability of an outcome is measured by the variance. As such, using information about both the magnitude of the variance relative to the expectations (to update learning rates) and the sign of the error (to update future choices), both signed and unsigned RiPEs may be crucial in an uncertain environment (Rushworth and Behrens, 2008). In our experiment, we found the unsigned RiPE appeared before the signed RiPE which may be similar to a two-component response in reward processing that has been shown before with the magnitude or surprise being encoded first followed by the valence of the stimulus (Schultz, 2016b). Future studies should determine if the unsigned RiPE is computed before the signed RiPE, or if the two occur almost simultaneously, although driving separate learning functions. Recent accounts on signed and unsigned reward prediction errors suggest that they may have different functions for learning depending on the time they occur during a trial (after the cue or after the feedback, Rouhani and Niv, 2021). Hence, the involvement of RiPEs learning can be studied with designs that separate them into post-cue and post-feedback errors. According to current evidence, they may have different effects on learning under volatility and stochasticity (Behrens et al., 2007). Or in other words, different types of RiPE may drive learning rates relative to recent or distant past experiences, or relative to the rate of change of the environment.

One limitation of the current design is the passive observation of stimuli selection and the absence of choice leading to rewarding outcomes. We implemented this procedure aiming to evaluate risk processing in relation to computational models and to isolate risk perception as much as possible from signals related to the preparation of a motor response (Bortoletto et al., 2011; Hsu et al., 2009; Preuschoff et al., 2006). This led to small effect sizes, expressed in the low discriminator performance and large attrition of participants whose threshold did not reach 50% accuracy of the discriminator. We requested participants to find and select the rewards they won on every trial, which required attention, and we instructed them that the points they won would be converted into pounds, which aimed to increase motivation. However, weaker effects may be due to a few factors including the lack of active participation leading to the outcomes, the

low value of the payoffs, and the secondary association between the presented social stimuli and the underlying rewards. Although previous studies have made efforts to dissociate motivation and salience from reward and risk coding (Preuschoff et al., 2006), motivation and attention are important aspects of goal-directed behaviours influencing transient dopamine levels (Brown et al., 2020; Roesch and Olson, 2004; Salamone and Correa, 2002). Hence, it is possible that decreased motivation and attention in our previous tasks may have led to weaker signals of risk processes. Another explanation is that risk signals are generally lower than those of rewards, as found previously with animal electrophysiology (Schultz, 2010). Future studies should explore the risk processes in a similar way that reward PEs were studied in the context of pure decision variables, goal-directed value, and action values (Rushworth and Behrens, 2008). Given our single-trial analysis combined with the leave-one-out cross-validation and our secondary validation establishing the parametric effects of the components, we are confident that the components we report code the expected risk and the signed and unsigned RiPE signals. However, without active choice and goal-directed decision-making, combined with weaker risk signals in general, one can expect small effects, especially when measured on the scalp as in the case with EEG. We suggest that future studies, aiming to implement our methods of passive risk perception, should control for different levels of motivation by implementing higher levels of rewards in terms of individual-specific utilities.

To conclude, our results show that there are three distinguishable components in the EEG signal in terms of their spatial and temporal profile. Moreover, consistent with non-social, probabilistic studies, we show that risk and RiPEs are coded at different time points after the onset of social stimuli, with specific spatial projections similar to previous non-social studies. Therefore, these results point toward a common system that is responsible for components of both social and non-social decision-making under uncertainty.

Risk-Seeking Driven by Stimulus Prediction Errors: Modelling and EEG Results

In the previous chapter, we showed spatio-temporal EEG patterns of risk-coding with face stimuli. We focused on recording errors associated with risk calculations. Now, we turn our attention to another type of error that may occur during the course of a trial as risky decisions unfold.

Recent research has suggested that trial-by-trial risk-taking might be influenced by cue-locked prediction errors, which can introduce biases into decision-making, influencing the likelihood of engaging in more or less risky choices. To delve into this further, I will first present a theory based on the mechanisms underlying dopamine functions in the basal ganglia, which predicts risk-taking behaviour. Subsequently, I will report the results of an experiment that tested this theory, comparing it to alternative risk-seeking models. Notably, the experiment encompassed both a non-social and a social condition, yielding significant findings that warrant separate discussions. Within this chapter, I will outline the non-social (baseline) behaviour, modelling, and EEG results and in the following chapter 4, I will elaborate on the social domain.

3.1 Background

Decision-making and learning models often rely on hypothesised rigid estimates of the environment to explain predictions that organisms hold. However, rewards (or more generally outcomes) are usually drawn from distributions which impose variability. In recent years, theories of Distributional Reinforcement Learning have started to influence modern computational models of learning (Bellemare et al., 2017; Dabney et al., 2020; Lowet et al., 2020; Zhao et al., 2020), and this is not surprising, given that under uncertainty it is adaptive and energy efficient to form an estimate of a ‘cloud’ of rewards, rather than trying to take account of each individual case or only hope for ‘the middle estimate’. Distributions can be informative estimates for an uncertain environment because apart from an average, they provide information about the variance. The uncertainty of options can be defined by the variance of rewards (Platt and Huettel, 2008), and under different conditions animals and humans decide to go for more uncertain or safer options (Kacelnik and Bateson, 1996). Dopamine (DA) levels have been linked to seeking exploration and preferences for uncertainty. In terms of structures, risky decision-making has been linked to the prefrontal cortex (PFC), the striatum, the cingulate cortex, the orbitofrontal cortex, the nucleus accumbens, and the basolateral amygdala (Platt and Huettel, 2008). The amygdala has been associated with increased risk preference with higher activation especially when executive function is compromised (Brand et al., 2007). A dopaminergic circuit between the basolateral amygdala and the nucleus accumbens was found to be related to a preference for larger uncertain rewards, yet biases arising in the amygdala are modulated by top-down control from the prefrontal cortex (PFC) (Jenni et al., 2017). These findings point out the complex interaction between cortico-striatal structures to code risk, and cortico-limbic connections to drive risk-aversion or risk-seeking.

Mikhael and Bogacz (2016) put forward a model that implicates cortico-basal ganglia connections and tonic dopamine levels as the main causes of trial-by-trial risk-taking. Action values are associations between actions and rewards, and have been proposed to be coded by differences in the synaptic weights of a direct (Go) pathway and an indirect (No-Go) pathway, populated by D1 and D2 dopamine receptors, respectively (Gerfen and Young, 1988; Mink, 1996). Mikhael and Bogacz (2016) proposed that the same pathways can code the mean of reward distributions, with increased neuronal firing, as well as the variance. In their paper, they denoted the synaptic weights of D1-populated (direct) and D2-populated (indirect) pathways on a certain trial as $G(t)$ and $N(t)$, respectively. They further suggested that although the average action value is coded by the difference in $G(t)$ and $N(t)$ (as proportional to the EV of the action), the uncertainty is coded by the sum of $G(t)$ and $N(t)$.

Figure 3.1 presents a schematic of cortico-striatal Go and No-Go connections. Two stimuli predicting a range of rewards that are either more desirable on average (green bubble), or less desirable (red bubble) bias the Go and No-Go pathways, respectively.




Figure 3.1 Theoretical influence of situational stimulus PEs on cortico-striatal activity. Go and No-Go pathways are shown in green and red, respectively. As explained in Chapter 1, they are influenced by learned weights from previous experience (via blue pathways from SN). In this example, the No-go pathway is influenced by a group of stimuli that are undesirable (in the red bubble) while the Go-pathway is activated by the desirable group (green bubble). The stimulus-PE is a result of the presentation of the red or the green bubble, and influences the No-Go or Go pathways, respectively.

Let us consider a separate situation in which an organism is presented with two stimuli. One stimulus is risky in that it predicts a distribution of rewards that has a higher variance (or a standard deviation), compared to a second safe stimulus that predicts a distribution with a lower SD. The basal ganglia risk-taking model proposes that the risky stimulus will be chosen over the safe one if the choice is preceded by a positive reward prediction error (Moeller et al., 2021). This preceding PE, termed the stimulus-PE, is not the carried-over PE from a previous choice, but rather the difference between the average value of all stimuli, and the stimuli presented at the current instance/trial (Mikhael and Bogacz, 2016). For example, an animal does not expect to find food, however, two sources of food are found, one risky, and one safe. Simply the presence

of food has already increased expectations (as it is better than not finding food at all), and the stimulus-PE is positive, which in turn would lead to a choice of the risky option. In some sense, this can be considered as an ‘optimism’ bias in an uncertain environment.

The mechanism behind risk-taking due to high PEs is explained in detail by Mikhael and Bogacz (2016), and I will only summarise it briefly. An important prerequisite is the finding that tonic dopamine has an influence on the direct and indirect pathways in which an increase in dopamine leads to the excitation of D1 and inhibition of D2 pathways (and action initiation), while a decrease in dopamine leads to the opposite pattern (and action inhibition). This has been supported by studies with genetics, Parkinson’s disease patients, and voltammetric and electrophysiological studies (Dodd et al., 2005; Frank et al., 2007; Gilbertson et al., 2020; Salamone et al., 2022; Salamone et al., 2005).

Figure 3.2 represents two contexts, one system with high tonic dopamine levels (top panel), and another with low tonic dopamine (bottom panel), while the blue and orange pathways show two possible learned actions of lower risk/SD (blue) and higher risk/SD (orange). The synaptic weights are indicated by the thickness of the arrows. The two options (blue and orange) have the same mean rewards, so, the orange arrows are thicker as the higher variance also indicates a higher reward compared to the blue option. As the orange pathways have overall higher weights, a high tonic dopamine level (which has an additive and excitatory effect on G) leads to a preference and selection of the risky option (orange). While a decreased dopamine level (bottom panel) acts on the N pathways, in which higher dopamine influence means inhibition of the action. This translates as inhibition of N2 and selection of N1. Overall, via situation-dependent dopamine increase or decrease, the cortex can bias the Go and No-Go pathways via an increase or decrease in tonic dopamine.

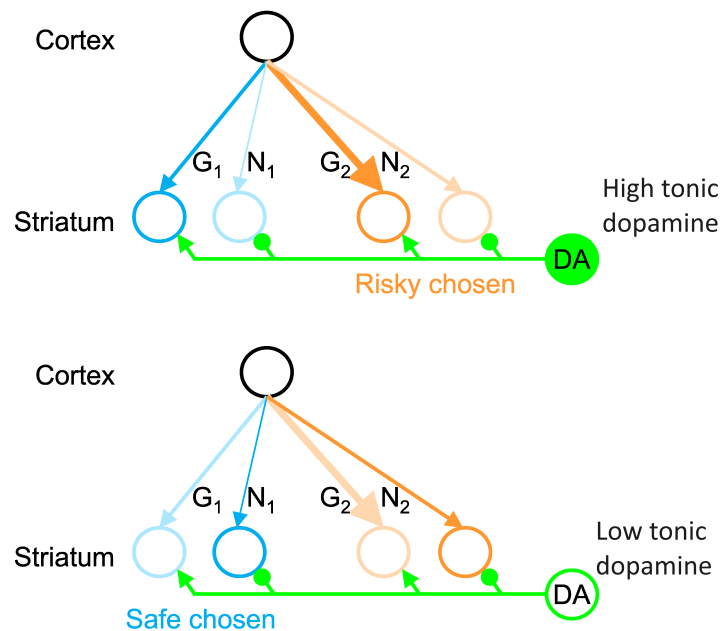


Figure 3.2 Dopamine influence on risk-taking. (Obtained from Mikhael and Bogacz, 2016) Two different actions are represented by either blue or orange paths and circles. The thickness of the cortico-striatal connections is indicated by the thickness of the arrows. Input from the cortex is received by the striatum via D1 (G_i) or D2 (N_i) connections. Tonic dopamine (green) either excites (arrows) or inhibits (dots) the D1 and D2 connections, respectively. The top and bottom panels show a system with high and low levels of dopamine, respectively.

The hypothesis that dopamine drives uncertainty has been supported in a few different ways. Genes controlling dopamine function in humans have been shown to predict exploratory behaviour separately in the basal ganglia and the PFC (Frank et al., 2009). Individuals with alleles predicting higher levels of dopamine showed increased exploration with the aim to improve learning (in the striatum) or (in the prefrontal cortex) ‘goal-directed’ exploration which is done in proportion to Bayesian uncertainty with the aim to improve current outcomes.

Another study showed that mesocortical dopamine can act on different networks to change reward-seeking and action-selection behaviours. Jenni et al. (2017) assessed rats and selectively disrupted D1 or D2-modulated PFC output pathways with dopamine antagonists. Disrupting both D1 and D2 led to a decrease in risky choices but with

different choice patterns. Disrupting D2 pathways (related to modulation of PFC on the amygdala) led to decreased risk-taking due to disrupted sensitivity to the frequency of rewarded and non-rewarded choices, while disrupting D1 reduced risky choices by increasing sensitivity to reward omissions. These results supported the idea that D1 pathways (PFC: NAc) promote risk-taking by reducing the impact of non-rewarded choices, and dopamine acting on D2 (but not D1) receptors promotes flexible patterns of reward-seeking via a network of PFC neurones that interfere with the basolateral amygdala. Although these results suggest that D1 and D2 pathways are both related to an increase in risk-taking, this was a consequence of long-term disruption of the pathways. It is unknown how trial-to-trial fluctuations in activity would affect risky behaviour. More importantly, the findings point to risky decision-making relying on midbrain structures such as the striatum and the amygdala, and top-down control via the PFC. So far, separate activation of D1 and D2 receptors on risk-taking has not been tested in humans, however, a recent behavioural and computational modelling study tested risk-taking in relation to this model (Moeller et al., 2021).

Moeller et al. (ibid.) tested the basal-ganglia learning model on risk-taking. The participants were presented with a two-arm bandit task and had to choose one of two images. As mentioned above, risk can be measured as the variance or standard deviation of the reward distribution. Each of the four images predicted rewards from a normal distribution with a pre-specified mean and SD. The results showed that after trial and error, the participants were able to learn and distinguish between the stimuli and to choose the ones with a higher EV when two images of different means were presented on the screen. When two stimuli with equal means but overall higher mean than the grand average were presented on the screen, participants were more likely to select the riskier option. Whereas when the two stimuli with an overall lower mean were presented, there was a higher probability of choosing the safer option. A model

explaining this behaviour via Stimulus PEs acting on risk-taking, termed the prediction error induced risk-seeking (PEIRS) was the best fit for the data and was able to predict risk-seeking behaviour better than eleven other models based on different reinforcement learning mechanisms.

The immediate influence of stimulus-PEs on risk-taking has not yet been studied in humans with neuroimaging. EEG is well suited to test a theory that makes predictions about subprocesses that may only be revealed within the span of a trial, namely, cue (stimulus) evaluation, feedback anticipation, and outcome assessment. In previous ERP studies, several components have been found directly after the presentation of a cue probabilistically predicting a reward, before the presentation of feedback, and post-feedback. The cue-P3 was found after cues predicting monetary gains and appeared 300 to 500 ms post-cue onset (Broyd et al., 2012). The same ERP was also stronger right after risky cues compared to ambiguous cues (Deng et al., 2023). These results suggest that the cue-P3 may be related to stimulus-PEs that are relevant for the saliency and predictability of the cues. However, no study has tested this ERP component as a potential driver of risk-seeking. Moreover, studies with ERPs focus on regions of interest, targeting a few electrodes that are selected based on previous empirical findings and averaging the ERPs across trials to deal with the low signal-to-noise ratio. These two methods can limit the potential to reveal spatial patterns across the brain and temporal dynamics across trials (Ouyang et al., 2017).

We implemented a model-based EEG linear discriminant analysis with the aim of testing the PEIRS model harnessing the high temporal resolution of EEG, which can allow the differentiation of prediction errors within the stages of a trial, and across the full task. We hypothesised that positive stimulus prediction errors, according to the model, will predict risk-seeking behaviour.

3.2 Methods

3.2.1 Participants

Participants were recruited through the Experiment pool of the School of Psychology and Neuroscience at the University of Glasgow. We performed a power analysis to determine the minimum sample size, using *pwr* package in R with a paired-samples t-test specification and the following: a significance level of 0.01, power of 0.8, and a *d* value of 0.71, obtained from the main t-test of Moeller et al. (2021). Based on this we estimated a minimum sample size of 27 participants. We collected 37 participants to account for exclusions based on too noisy EEG data due to muscle artifacts or equipment issues, and poor performance on the behavioural task. From the 37 participants, 4 were excluded from the behavioural and EEG analysis (2 of them due to poor performance on the behavioural task (Figure 3.3 - performance), and 2 due to technical issues leading to missing EEG data for half of the duration of the experiment). Poor behavioural performance was defined as how often a stimulus with a higher average is chosen between two stimuli with different means (see Task). The final sample was 33 participants (age: min = 19, max = 37, median = 26, sex: 16 females, 17 males).




Figure 3.3 Performance on the behavioural task across participants. Excluded are in grey.

3.2.2 Task

For a two-arm bandit task, we used 4 stimuli on each block, differing in colour and having a unique symbol at the top part (example stimuli - top of Figure 3.4). We aimed to design the stimuli to appear as ‘slot machines’ while minimising shading and complexity.

On each trial, the participants were presented with a fixation cross for a variable duration of 1 to 1.5 seconds (Figure 3.5). Then two out of the four possible options for slot machines were presented on the screen. Participants had 2 seconds to make a response on a Cedrus box with a left or a right button for the left or right slot machine according to their position on the screen. The lever on top of the chosen slot machine appeared as if it was pressed for a duration between 1 and 1.5 seconds, after which a reward in points was displayed inside two grey boxes of the chosen slot machine for 1 second. If no response was made within 2 seconds of the stimuli being shown, the message ‘Too slow!’ appeared on the screen and the next trial started. The main task consisted of 3 blocks, each lasting approximately 15 minutes, and including 180 trials (total: 540 trials per participant). Each block consisted of slot machines associated with four different distributions. Two of the machines returned 60 points on average (both-high pair), while the other two returned 40 (both-low pair). The both-high and both-low pairs each consisted of a risky option (standard deviation of 20) and a safer option (standard deviation of 7).

The slot machines changed on each block. Although the four colours remained the same across blocks, the symbols inside each slot machine changed, indicating that the machines are different, meaning the participant must re-learn the associated values again. The participants were made aware that they may develop a preference for some machines compared to others over the course of a block, and that the machines will

change symbols and associated rewards when a new block starts. So, participants were also told that they had to re-learn the associated values via trial and error separately for every block. During each block there were three breaks of 30 seconds each, on which the total amount of points won up to that moment was shown.

The design of the experiment was a within-subjects design with respect to the social conditions and stimulus prediction errors (slot machine distributions). The main task always started with the baseline in which participants were not observed, while the last two blocks were randomised in terms of the social conditions (detailed in Chapter 4).

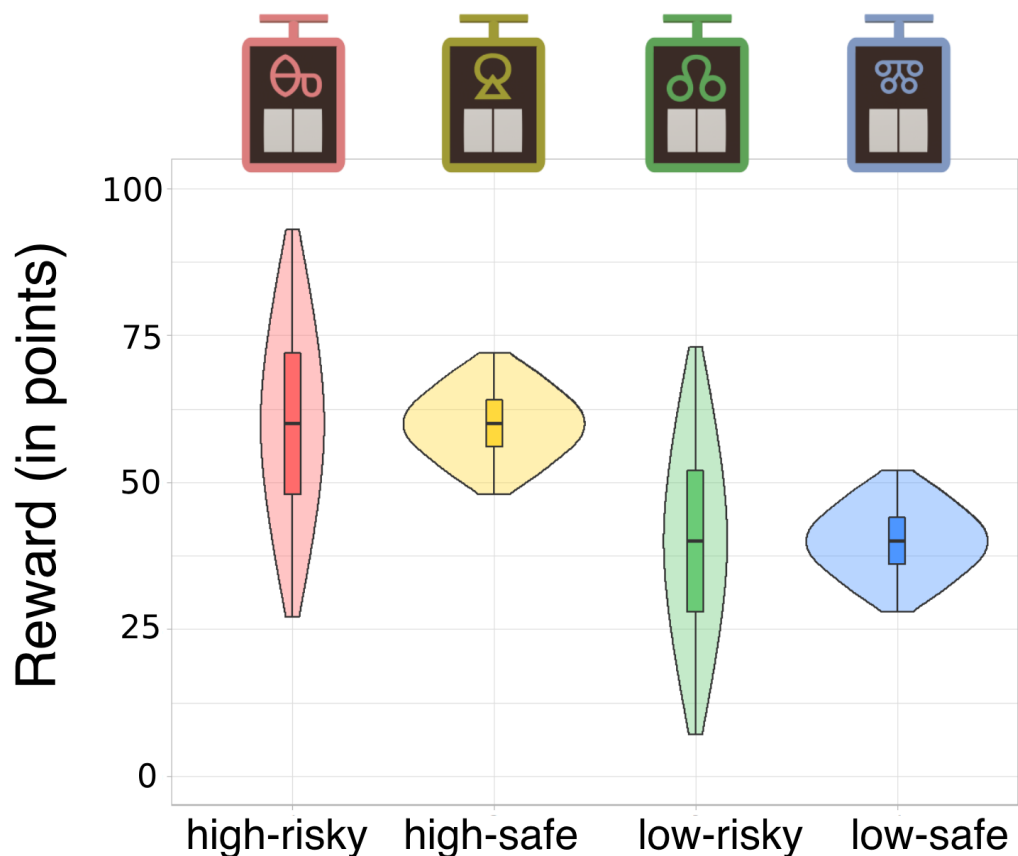


Figure 3.4 Four example slot machines with their associated reward distributions. 'High' refers to a high mean (60 points), and 'low' is the low mean (40). Risky refers to the SD of 20, and safe refers to the SD of 7.

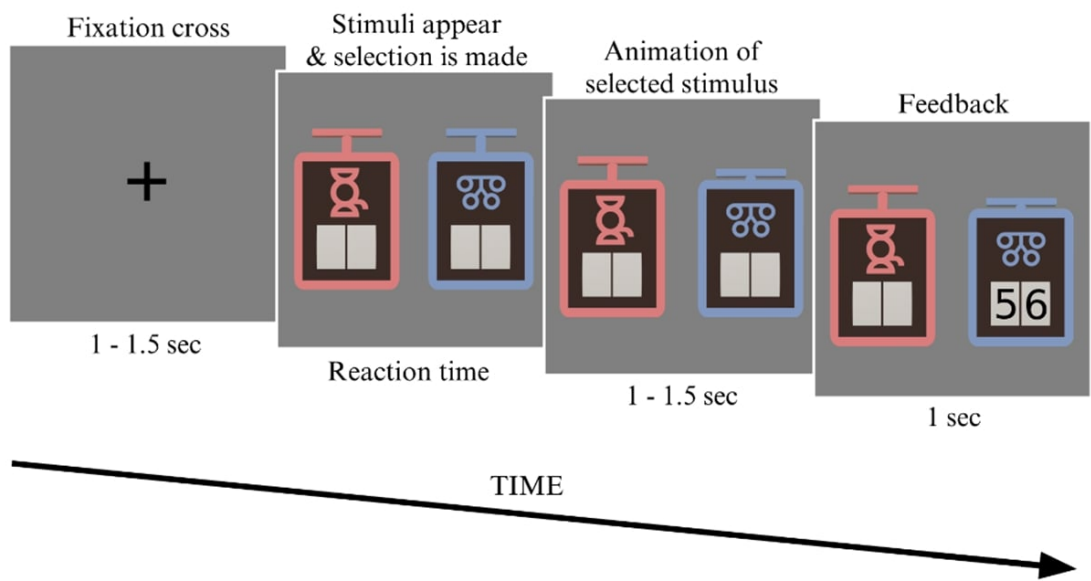


Figure 3.5 Timeline of the bandit task. The first screen was a fixation cross with a variable duration. Then the participants were presented with a pair of slot machines for up to 2 seconds, or until a choice was made. A choice triggered an animation in which the lever of the chosen stimulus appeared as if pressed down for 1 to 1.5 seconds. Finally, the feedback appeared in points for 1 second as two digits within the grey boxes on the chosen stimulus. The next trial started with a fixation cross.

3.2.3 Theoretical analysis

For the theoretical analysis and model fitting we used the theory outlined by Mikhael and Bogacz (2016) and a toolbox in Matlab Moeller et al. (2021). To summarise the theory briefly, it proposes separate action channels for each option of slot machines in the task, and the probability of choosing each option is dependent on the activation of the relevant action channel. Figure 3.5 shows a schematic of the process for both-high (A) and both-low (B) combinations of stimuli. The activation of each action channel is the difference between Go (G) and No-Go (N) channels activations, innervated by D1 and D2 dopamine receptors, respectively. In line with the Rescorla-Wagner learning rules, Q is the mean of the reward distribution, while S , is the spread (SD). The model further proposes that the spread of each distribution influences the activation of action channels via dopamine. After learning the association between the stimuli and their distributions, during the fixation cross dopamine is proportional to Q (the mean of all rewards). If dopamine levels are increased (by a positive PE) when two stimuli with higher (than the grand average) mean are presented, S contributes positively to action activation, hypothetically leading to the stimulus with a higher spread being chosen. On the other hand, if PEs are negative compared to baseline as in the case when the two stimuli with lower means are shown, S reduces action activation, leading to a biased choice towards the safe stimulus.

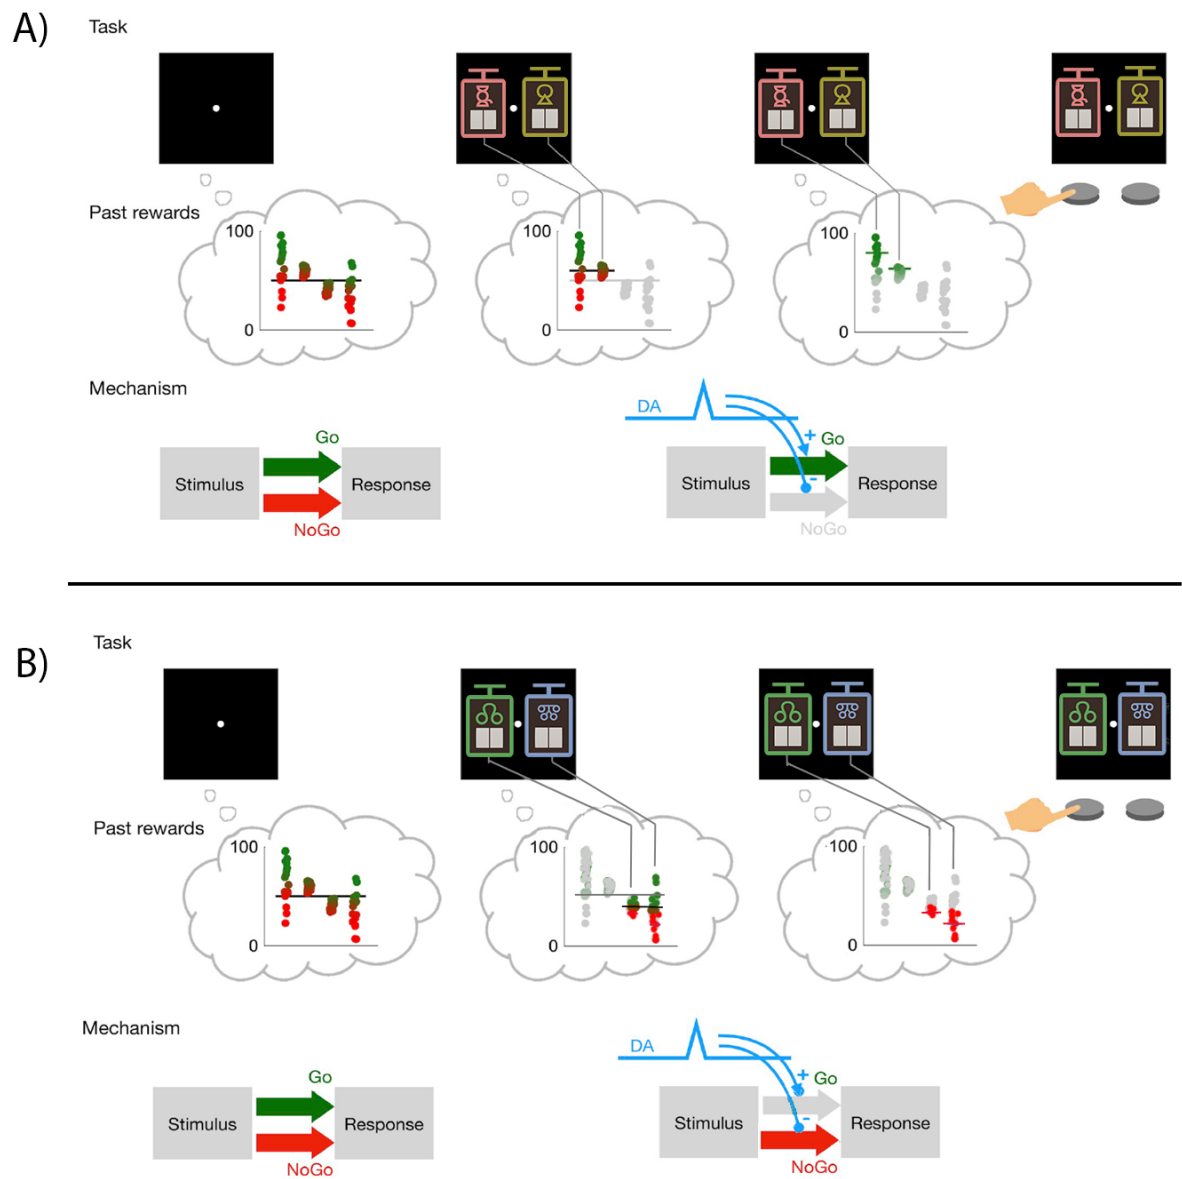


Figure 3.6 A schematic of two instances of slot-machine combinations and the proposed basal ganglia mechanism for each one (modified from Moeller et al., 2021). As the task progresses there is a different representation of the possible outcome values. At the beginning of the task, all values are expected. With the presentation of the slot-machine pair, expectations are updated to the relevant values. In case A) the slot machines have a higher combined average than the overall mean leading to an increase in dopamine that activates the Go pathway. This in turn leads to the selection of the stimulus with now higher local average (green dots on the third bubble of A). In the case of B), the two slot machines have a lower average than the grand average of all stimuli. The dopamine function activates the NoGo pathway and now the focus is on the lower options (red dots on the third bubble of B). This leads to the selection of the option with a higher local mean, which happens to be the safer option.

The different weights put on action values described above are influenced by the stimulus-PE and captured by a model developed by Moeller et al. (2021), PEIRS (for Prediction Errors Induce Risk Seeking). Within this model the stimulus PE is defined as:

$$\delta_{stimulus} = \text{mean}_{shown}(Q) - \text{mean}_{all}(Q) \quad (3.1)$$

Hence, the stimulus PE (stim-PE) would be +10 (60 minus 50) in the case of two slot machines with higher means, and -10 (50 minus 60) if the two slot machines with lower means are presented. The outcome prediction error is the difference between the mean of the selected stimulus and the outcome number of points presented at the feedback screen ranging from 1 to 100:

$$\delta_{outcome} = r - Q_{selected} \quad (3.2)$$

With r being the outcome in points. From then on, the activation of a channel is defined as:

$$A_i = Q_i + \tanh(\omega\delta) \times S_i \quad (3.3)$$

Where A_i is the activation of channel i . The term $\tanh(\omega\delta) \times S_i$ modulates the impact of the prediction error on the action value. The hyperbolic tangent function compresses the prediction error between -1 and 1 to represent plausible neural mechanisms (due to its sigmoid shape, saturation properties, and differentiability, Abbott and Dayan, 1999; Maass, 1997). Multiplying it by S allows for incorporating the uncertainty of

the reward distribution. ω is a rescaled proportionality constant that determines the sensitivity to prediction errors. The resulting value is added to the mean reward value Q estimated at the current trial. As such, the activation of a channel is influenced by the prediction error and the spread of the distribution.

We used a modified PEIRS model, *PEIRS2*, in which we added an additional free parameter, representing a bias on the stimulus-PE:

$$A_i = Q_i + \tanh(\omega(\delta + \tau) \times S_i) \quad (3.4)$$

Where τ is a free parameter and $\tau > 0$. This modification expands the flexibility of the model, enabling the incorporation of an additional factor that can modulate the influence of the prediction error. This factor can be individual differences or context-dependent effects and is motivated by the fact that our study differed slightly from the original paper that incorporated PEIRS (Moeller et al., 2021). (Relevant for the following chapter, this parameter allows the model to be generalised across our social and non-social conditions and to capture more variance in the estimation of the action values).

The models we tested are all based on the RW model. Although we tested all models provided by the toolbox and explained in detail in Moeller et al. (ibid.), we report six models that explained the data best, with the rest of the models performing poorly in terms of BIC and comparisons between the simulated and obtained data.

To test risk-seeking tendency as caused by the nonlinearity of a utility function, and more specifically a utility function with differences between wins and losses as explained by prospect theory, we used the *s-shaped UTIL* (Trepel et al., 2005). In our case, wins are treated as positive outcome PEs, and losses are treated as negative outcome PEs relative to a subjective reference point. This theory proposes concave curvature of the function for positive PEs, and a convex function for negative PEs.

$$z = \frac{r - m}{\sigma} \quad (3.5)$$

$$U = m + \sigma \times \text{sign}(z) |z|^{k_{s\text{-shaped}}} \quad (3.6)$$

In the above equations, rewards are denoted by r , while m represents the certainty equivalent, or baseline level of utility when there is no uncertainty or risk involved, and σ is a scaling factor that represents the degree of risk aversion. This parameter determines the degree of curvature in the value function. In prospect theory, the value function is typically concave for gains (meaning that people are risk-averse when it comes to gains) and convex for losses (meaning that people are risk-seeking when it comes to losses). The parameter σ controls the shape of this value function. A larger σ makes the value function flatter and, thus, makes people more risk-averse. A smaller σ makes the value function steeper and, thus, makes people more risk-seeking.

Z is a random variable, representing the deviation from the expected outcome (providing its sign as $\text{sign}(z)$, and magnitude, $|z|$). $k_{s\text{-shaped}}$ is a parameter that controls the steepness of the curve, hence, the tolerance to extreme outcomes. The s-shaped UTIL model has free parameters $\alpha \in [0, 1]$, $\beta > 0$, $k_{s\text{-shaped}} \in [0, 1]$ and fixed parameters $Q_0 = 50$, $\sigma = 50$, $m = 50$.

Another model we tested was the scaled PE, which tested the hypothesis that PEs adapt to the statistics of the reward distributions, changing according to the current context. This model scales PEs relative to the perceived reward variability, which is tracked as well as the average reward. Risk-related biases would arise due to stimulus-specific learning speeds that are influenced by the scaling.

$$\tilde{\delta}_{outcome} = \frac{r - Q_{chosen}}{S_{chosen}} \quad (3.7)$$

$$\Delta Q_{chosen} = \alpha_Q \times \tilde{\delta}_{outcome} \quad (3.8)$$

$$\Delta Q_{chosen} = \alpha \times |\delta_{outcome}|^{k_{attention}} \times \delta_{outcome} \quad (3.9)$$

The model has free parameters $\alpha_Q \in [0, 1]$, $\alpha_S \in [0, 1]$, $S_0 > 0$, $\beta > 0$, and a fixed parameter $Q_0 = 50$.

With yet another model, *Attention RATES*, we tested for learning rates that update according to the absolute PE on each trial, with higher PEs in terms of magnitude contributing more to attention and potentially to memory.

$$\Delta Q_{chosen} = \alpha \times \left(|\tilde{\delta}_{outcome}| - 1 \right) \quad (3.10)$$

This model has free parameters $\alpha \in [0, 1]$, $\beta > 0$, $k_{attention}$ and a fixed parameter $Q_0 = 50$.

Finally, with *OEIRS* (Outcome errors induce risk seeking) we tested for the possibility that outcome PEs on the previous trial could drive risk-seeking in the current trial. This model is identical to the PEIRS model, however, in equation Eq. 3.3 the $\delta_{outcome}$ (as in Eq.3.2) is used instead of $\delta_{stimulus}$.

3.2.4 Additional materials

Materials for the EEG are the same as our previous experiment, detailed in Chapter 2 - EEG data acquisition. Additionally, we recorded three questionnaires at the end of the experiment. One questionnaire was about general feelings about the in-group and out-group observers filled out before the debriefing was done. A questionnaire combining the Risk Propensity Scale and the Brief Sensation Seeking Scale (both denoted as RPS) was administered (Hoyle et al., 2002; Meertens and Lion, 2008). A composite score was obtained by first reverse scoring the appropriate items and then obtaining the sum of all scores per participant and turning this into a percentage of the maximum possible score. And the third measure was STAI, a measure of state and trait anxiety (Charles et al., 1972). The STAI answers were forward and reverse-scored according to the convention (ibid.) and added together to form one score for SAI and one for TAI.

3.2.5 Procedure

Each participant completed an online example (via Pavlovia.org) of the main task one day before the experiment. Once in the lab, participants were given more information about the experiment and signed a consent form. Moreover, the participants were told that they could earn any amount in the range of £12 to £18 depending on the number of points they won. Due to the convergence to the average, independent of choice, the final points won were very similar across participants, hence, they were paid £16 or

£18 if the EEG set-up took longer than expected. Participants were instructed that the task consisted of three blocks, one in which they would the task on their own, followed by two blocks during which another participant would observe their performance. The details of the social procedure are reported in Chapter 4.

The only information given about the behavioural task was that the aim was to choose one out of any two slot machines paired on the screen to maximise the number of points won throughout each block. As the participants performed the task online before coming to the lab (with different slot machines), no further training was provided.

The EEG cap was then applied, and the participants were assigned to a social condition (Chapter 4). Participants were invited into the shielded EEG booth where impedances were reduced to less than 20 kOhm. Participants completed an eye calibration task for later EEG eye movement and blink removal as well as another eye calibration for the eye tracker. After the main task was completed, the participants filled out additional questionnaires and were paid and debriefed.

3.2.6 Data analysis

The behavioural data were analysed in R (v.4.2.3) and Matlab (v.R2022a). (Some of the results are reported using the *report* package in R, Makowski et al., 2023). The behavioural data was obtained by dividing the trials into those when two stimuli with an overall higher mean were shown (both-high), those with different means (different), and those with an overall lower mean (both-low). The interest was mainly on the proportion of risky choices during both-high compared to the risky choices made during the both-low condition. The theory proposed that the former would show more risk-seeking, while the latter – less risk-seeking. To be able to closely compare our results to those of Moeller et al., 2021, we also followed their analysis pipeline. We performed a t-test on the average probability of risky choices during both-high compared to both-

low conditions ($p(\text{risky}|\text{both-high})$ vs $p(\text{risky}|\text{both-low})$). As there are two values per participants, hence, within-subjects comparisons, we ran a paired-samples t-test. For every t-test we made sure that the assumptions of normality and homoscedasticity were met by running Shapiro-Wilk tests on each level of the variable 'pair type'.

As there is some learning throughout the experiment, we also divided all 180 trials into 5 and computed the same proportions of risky compared to safe choices at each timepoint for both-high and both-low conditions. We plotted these results as a learning curve. The reaction times during the three conditions were positively skewed, so a mixed-effects model was used (Lo and Andrews, 2015) with participant ID as a random factor, pair type (both-high/both-low/different) as a predictor, and reaction times in milliseconds as an outcome. This was done using the *afex* package in R (Singmann et al., 2016).

To fit the computational models and obtain the simulations we used the VBA toolbox (Daunizeau, n.d.) and code provided alongside Moeller et al. (2021). This toolbox employs a Variational Bayes approach to analyse a set of measurements. By inputting a generative probabilistic model, and prior distributions over model parameters, the toolbox provides outputs including an approximate posterior distribution over model parameters, an approximate posterior distribution over latent variables (unobserved variables), and fit statistics such as the Bayesian Information Criterion (BIC). The mean of the posterior distributions obtained from the toolbox outputs was used to estimate the model parameters and latent variables, including values and prediction errors. Simulations were performed using parameters from the fits. Each model produced 33 parameter sets corresponding to the 33 participants in our study. Using these parameter sets we simulated our task and generated datasets of the same size as the original dataset. This process was repeated 1000 times to ensure the stability of the resulting distributions. From the simulated datasets, we extracted 33000 pairs of aggregate risk preferences and risk traces (risk preferences as a function of trial number) for each model.

MATLAB's `ksdensity` function, a kernel smoothing function, was used to estimate the underlying probability density of the distribution of risk preferences. The estimated probability densities were then visualized using MATLAB's `contour` function.

Similar to the behavioural analysis, we first segregated the risk preferences based on different conditions. Next, we calculated the average risk traces by averaging the traces over repetitions of the experiment for each simulated participant. Finally, we further averaged these traces across participants. To enhance the display, the averaged simulated traces were smoothed using a 20-point moving average filter.

We also computed the average difference in risk preference between conditions for each simulated participant in each simulated experiment, resulting in 33000 differences in risk preferences for each model. These differences were then averaged across experiments to obtain 33 differences in risk preferences per model. To compare these distributions, we considered the empirical distribution of mean differences in risk preference between conditions, which also consisted of 33 data points (one for each participant).

The EEG data were pre-processed with an in-house pipeline and analysed with a single-trial multivariate analysis (pre-processing is explained in detail in Chapter 2 EEG data pre-processing). The equations and process for the single-trial LDA are shown in Chapter 2 Single-trial EEG analysis (equations 2.4 and 2.5).

3.3 Results

The aim of this experiment was to test the hypothesis that positive stimulus-PEs drive risk-seeking and to shed light on the spatiotemporal pattern of the stimulus-PE and its relationship with risky behaviour. All of the following results are for the baseline condition only which was always completed by the participants before the two social conditions.

3.3.1 Behavioural results

We proposed a higher proportion of risky choices in the both-high compared to both-low trials. We plotted this as a scatter-plot (Figure 3.7) between $p(\text{risky}|\text{both-high})$ vs $p(\text{risky}|\text{both-low})$. If the results are in line with the theory, we would expect more points to be towards higher $p(\text{risky}|\text{both-high})$ and lower $p(\text{risky}|\text{both-low})$, hence, more points in the low right triangle of the plot.

As shown in Figure 3.7, we found that participants were more likely to select the riskier slot machine in the case when the presented slot machines were with a higher EV of 60 points. Whereas if both slot machines were with lower EVs, the safer one was selected more often. This was confirmed by a paired-samples t-test, $t(32) = 3.35$, $p = 0.014$, mean difference = 0.12, 95% CI [0.03, 0.22]. 63% of participants showed a higher probability of risk-taking during both-high compared to both-low condition.

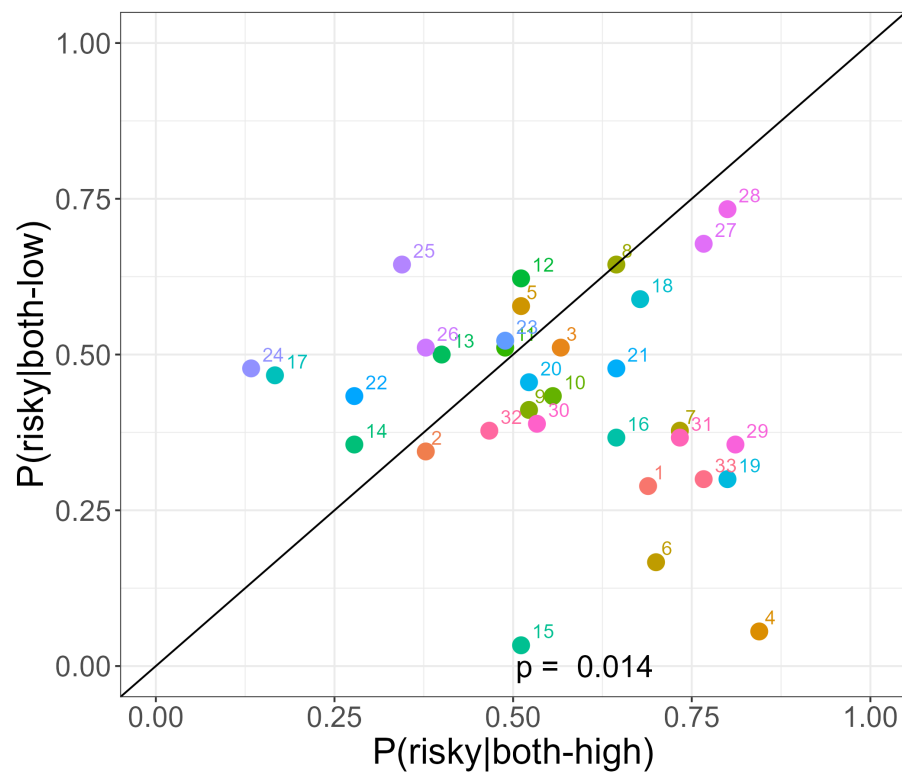


Figure 3.7 Risky choices for both-high and both-low options. Each point is a participant, $N = 33$. According to the theory, one would expect more points to be in the low right triangle.

The above tendency was further shown in the learning curve across trials in Figure 3.8 below. As trials progressed, participants were more likely to select the risky option in the both-high condition, compared to the both-low. Figure 3.8 reveals the average pattern of risk-taking across the span of the full task, each bin consists of 36 trials. An ANOVA was performed to determine the interaction between binned trials and pair type (both-low/both-high). The stars indicate the instances of significant differences between both-high and both-low trials in $p(\text{risk taking})$ corrected for multiple comparisons (Tukey HSD).

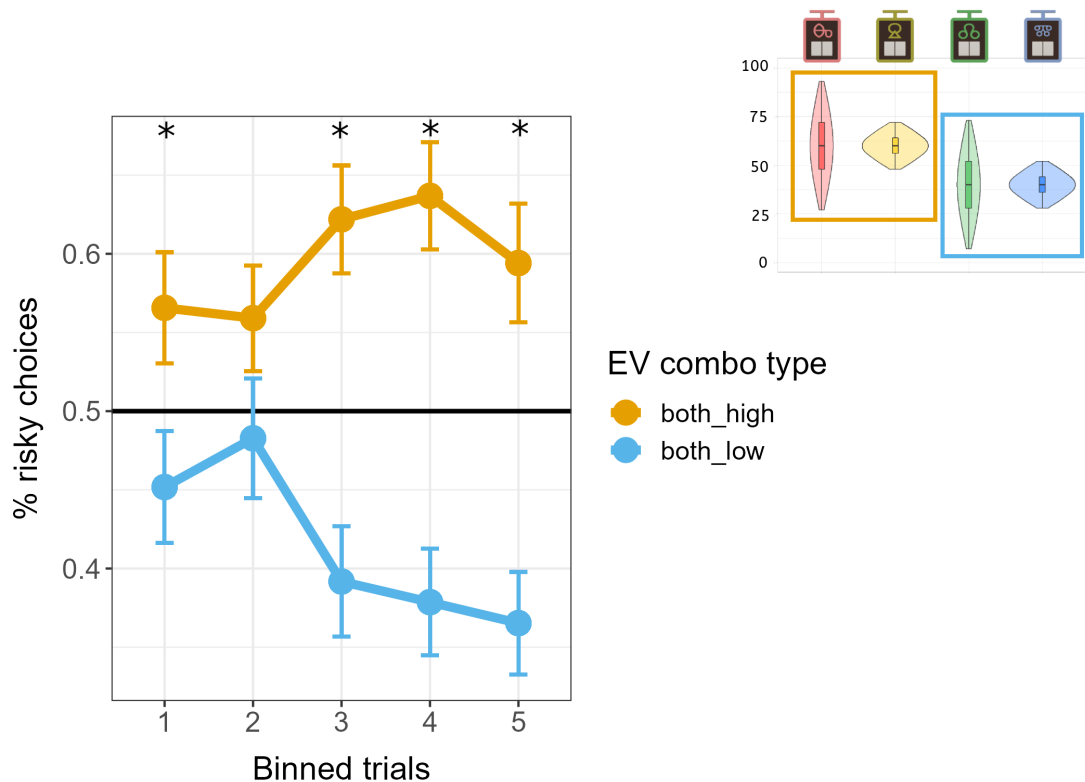


Figure 3.8 Learning curve for risk-taking during both-high and both-low combinations. Shown across the task (180 trials), each bin contains 36 trials. On the right is a schematic reminder of the pairs of both-high EV (orange frame) and both-low EV (blue frame), with risky distributions in red and green, and safe in yellow and blue.

Further, reaction times were faster when stimuli with both-high means or different means were shown compared to both-low means. This is shown in Figure 3.9, and confirmed a linear mixed-effects model. Using both-high pair as reference level: the difference with both-low pair was statistically significant and positive (beta = 0.22,

95% CI [0.16, 0.28], $p < .001$) and the difference with different-means pair was non-significant and negative (beta = -0.02, 95% CI [-0.07, 0.02], $p = 0.266$), re-referencing to both-low, the difference with different-means pair was statistically significant and negative (beta = -0.24, 95% CI [-0.29, -0.20], $p < .001$).

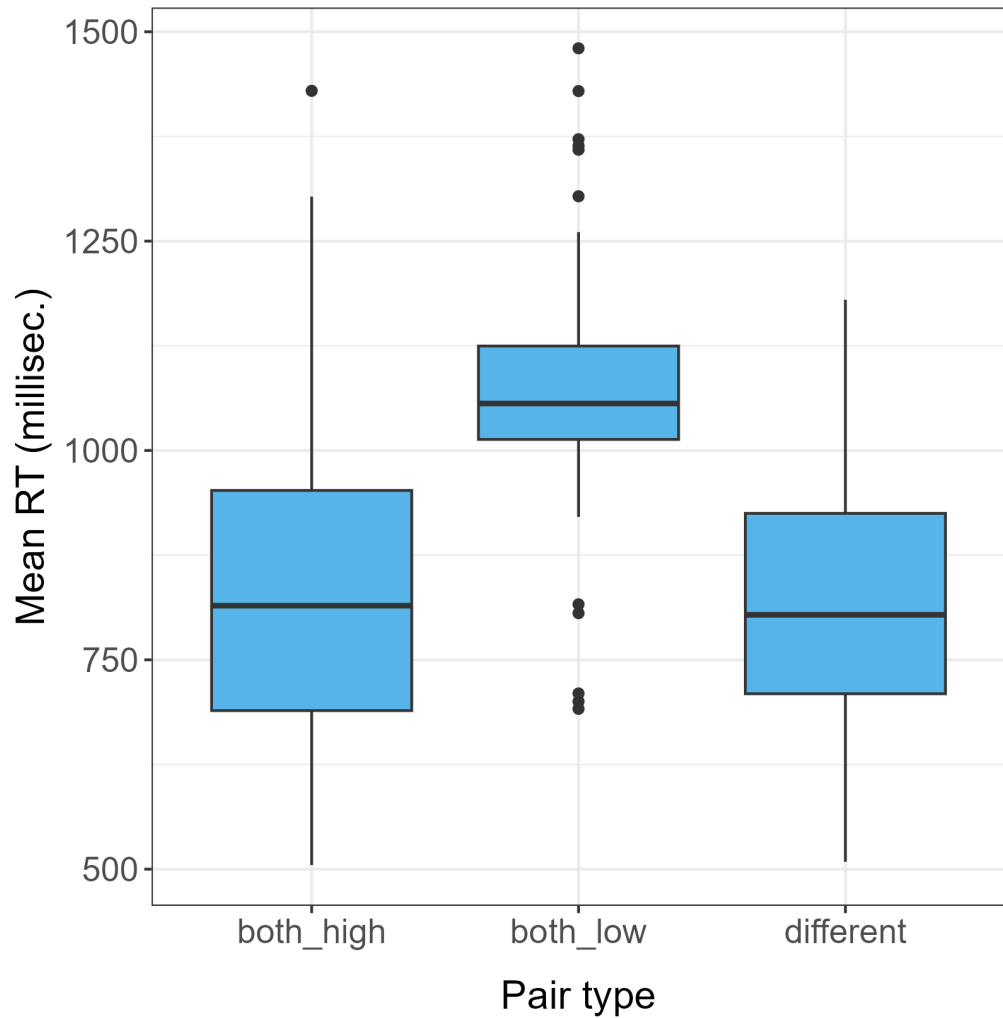


Figure 3.9 Reaction times to choose an option when presented with each pair of stimuli.

Additionally, we tested a correlation between risk-taking and the two questionnaires (one for general risk-taking propensity, and another for state and trait anxiety) using Pearson's r . We found no significant correlation between the average score of participants on each questionnaire and any of the risk-taking measures in the task (overall risk-taking, risk-taking during both-high, and risk-taking during both-low), as shown on Figure 3.10.




Figure 3.10 Pairs plot between risk-taking for each condition, anxiety scales, and risk-propensity scale. The variables are: percent risky choices for both-high, both-low, different-means conditions; SAI scores, TAI scores, and risk propensity scale score.

3.3.2 Model comparisons

Figures 3.11 and 3.12 show the results of the recorded behavioural data and simulation results using the obtained parameters for each model. Fig. 3.11 A) shows the BIC values for each model summed over all participants. A lower BIC value indicates a better fit. Panel B) shows the differences in risk preferences across the two conditions of interest (both-high and both-low). Blue bars indicate the simulated datasets with parameters obtained from each model fit (with SE error bars). The grey bar shows the original experiment results. The models that produce the closest results to the data are indicated as not significantly different (n.s.) from the original risk preferences. This was obtained using separate t-tests. Figure 3.12 A) shows the population risk preferences in simulated and experimental datasets across the two conditions. Blue contour plots show the distributions extracted from simulated data with experimental data superimposed as a grey scatterplot. The shading corresponds to estimated probability density functions and is scaled to match each distribution's density value range. Here we considered the best models as the ones that most closely matched the distribution of data points. B) shows the learning curves of risk preferences for each condition. Grey and black lines show the original dataset, while darker and brighter blue shows the simulated datasets with the shaded area being the SE across simulated participants (smoothed using a moving average). For this visualisation, we considered as best the models that resulted in the most similar traces to the original dataset, namely, both-high and both-low slowly diverging from baseline in opposite directions.

As shown in Figure 3.11 B, four models resulted in similar differences between the proportion of risky choices in the both-high and both-low conditions. Out of the four, only three fit the data across trials well, PEIRS, PEIRS2, and s-shaped UTIL as shown on the same figure, panel C. From these three, the best two to fit the experimental data distribution on panel D) were PEIRS and PEIRS2, which did not produce the narrow distributions observed for the other models. However, the model with the lowest BIC was PEIRS2.

The PEIRS2 model had an additional free parameter that could account for individual differences or contextual differences as an additional bias on the interpretation of the stimulus-PEs in particular. The BIC is a measure used for model selection, balancing the trade-off between model complexity and goodness of fit, and free parameters can have either a positive or a negative effect on this metric (Király and Hangya, 2022). On one hand, increasing the complexity of a model with more free parameters can potentially improve the model's ability to capture intricate patterns and fit the data more accurately. A more complex model may achieve a better fit and minimize the residual errors, resulting in a lower BIC value. However, BIC penalizes model complexity by including a penalty term that grows with the number of free parameters as these may lead to overfitting the data. So, if two models are identical apart from a free parameter, the model with a free parameter will be penalised and result in higher BIC. In our case, however, PEIRS2 has an additional parameter, however, remaining with the lowest BIC (of 2536) compared to PEIRS (BIC = 2546), hence, it strikes a balance between complexity and performance. Therefore, we subsequently used the parameters from PEIRS2 to extract stimulus and outcome prediction errors across trials to be used in the EEG analysis.

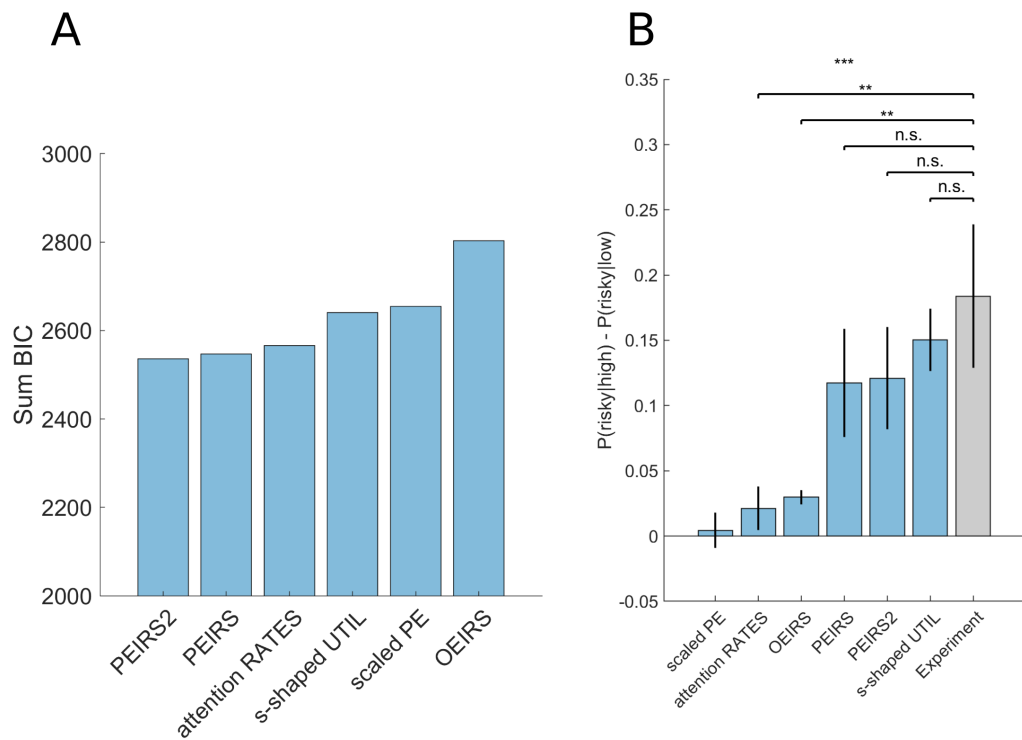


Figure 3.11 Model comparisons for the baseline condition. BIC and model-simulated data. A) BIC of each model, with lower BIC indicating a better fit. B) Average probability of choosing risky options in both-low condition subtracted from the probability of choosing risky during both-high condition.

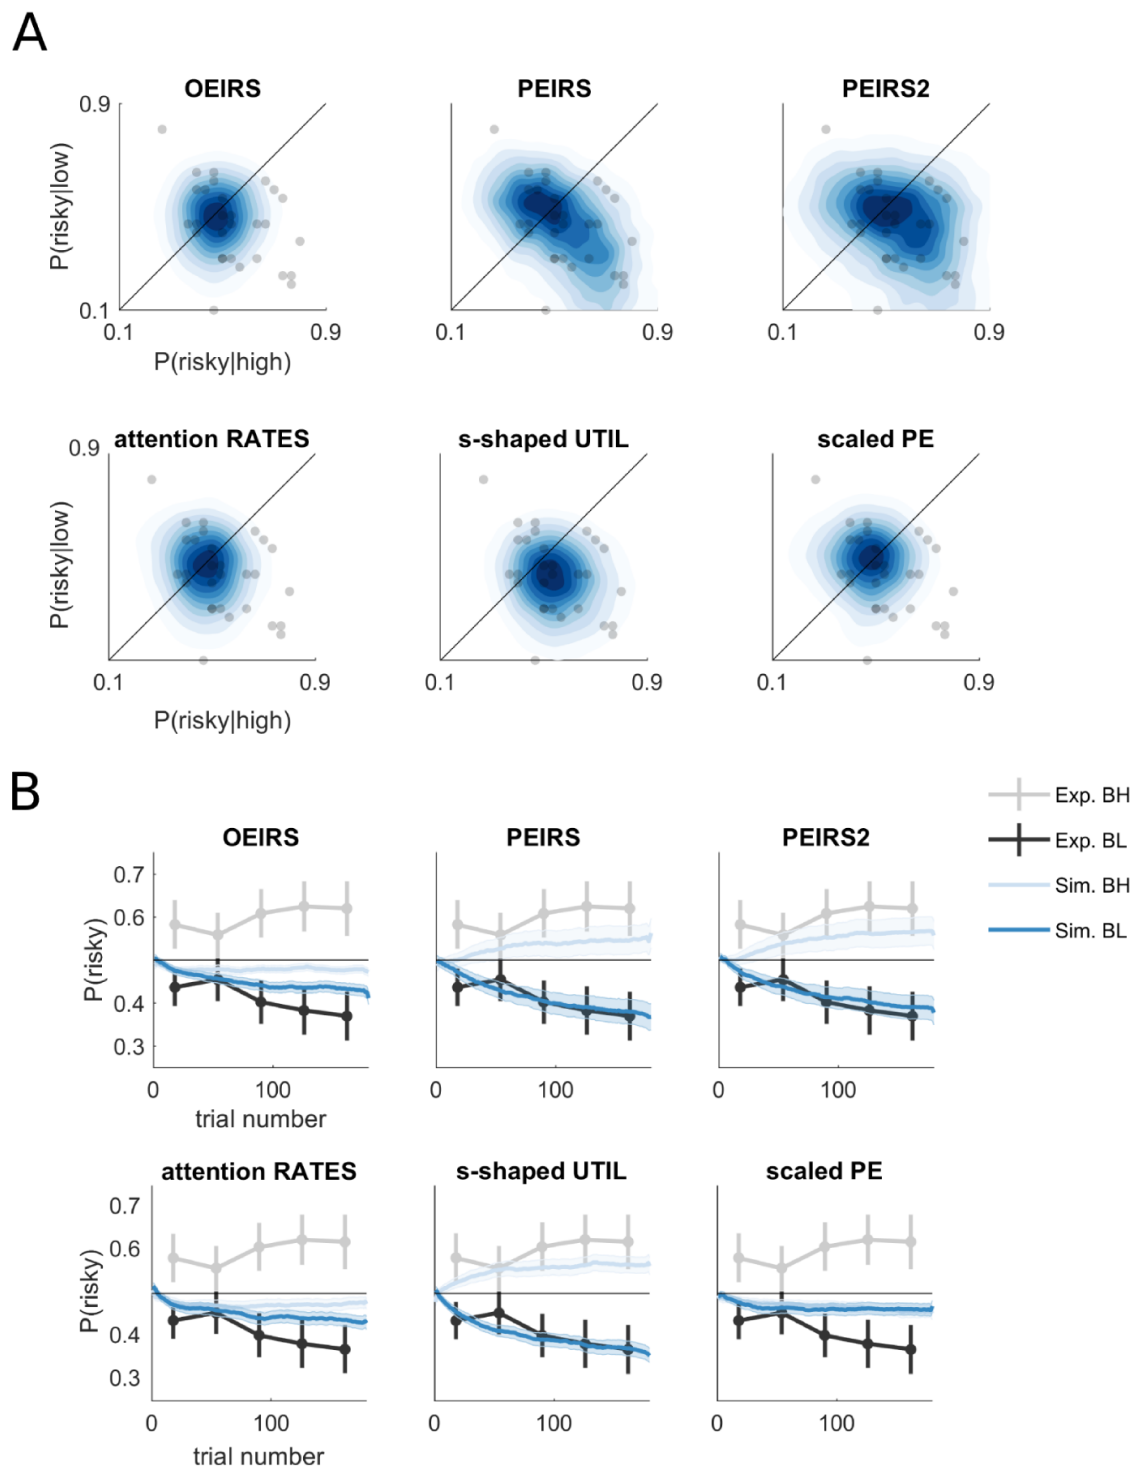


Figure 3.12 Model comparisons for the baseline condition - 2. A) Risk preference distributions from simulated datasets are shown as blue contour plots (shading corresponds to estimated probability functions). Superimposed experimental data per participant is shown as grey dots. B) Experiment data across trials shown with grey (Experiment both-high - Exp. BH) and black (Experiment both-low - Exp. BL) lines and error bars (SE), and simulated data for both-high are shown with bright blue lines (Sim. BH), while simulated data for both-low are indicated with dark blue (Sim. BL).

3.3.3 PEIRS2-stimulus-PEs predicting risk-taking

The PEIRS2-obtained stimulus-PEs significantly predicted risk-taking on a trial-by-trial basis. We fitted a logistic mixed model (estimated using ML and Nelder-Mead optimizer) to predict risky choices with PEIRS2 stimulus-PEs (formula: $\text{risky} \sim \text{PEIRS2-stim-PE}$). The model included ID as a random effect (formula: $\sim 1 \mid \text{ID}$). The model's total explanatory power was weak (conditional $R^2 = 0.05$), and the part related to the fixed effects alone (marginal R^2) was of 0.04. The effect of PEIRS2 stimulus-PEs was statistically significant and positive (beta = 0.05, 95% CI [0.04, 0.06], $p < 0.001$; Std. beta = 0.36, 95% CI [0.30, 0.43]).

3.3.4 Model-based EEG results. Signed Stimulus-PE

The baseline data of each participant was divided into model-based negative, medium, and positive stimulus-PEs. These levels were obtained by arranging all PEs and dividing them into three equal parts of 60 trials each (a total of 180 trials). The LDA was then trained with the EEG data from trials with negative and positive stimulus-PEs, while it was tested on the middle-PE trials (mostly corresponding to trials when the two stimuli with different means were presented). This resulted in a discriminator performance that shows the maximum discrimination of the EEG space between trials with negative and positive stimulus-PEs (which fell closely within both-high EV and both-low EV trials respectively). The discriminator result was validated with a leave-one-out procedure, and an area under a receiver operator curve was computed at every timepoint around the stimulus onset (termed Az-values). The result is considered at a chance level if at 50% (0.5 Az) or below.

The single-trial multivariate analysis successfully discriminated between the trials with negative versus positive stimulus-PEs as shown in Figure 3.13, and indicated by Az-values higher than 0.5. The discriminator peaked at 500 ms after the onset of the two slot machines. Reaction times showed that participants responded on average 1 sec after the onset of the stimuli. The scalp topographies shown in Figure 3.13 are obtained from the average of 33 participants first between 150 ms to 250 ms and then from 450 ms to 550 ms. For the remaining analysis, we focused on the second window due to its higher discrimination and larger number of consecutive time windows above 0.5 Az across participants.

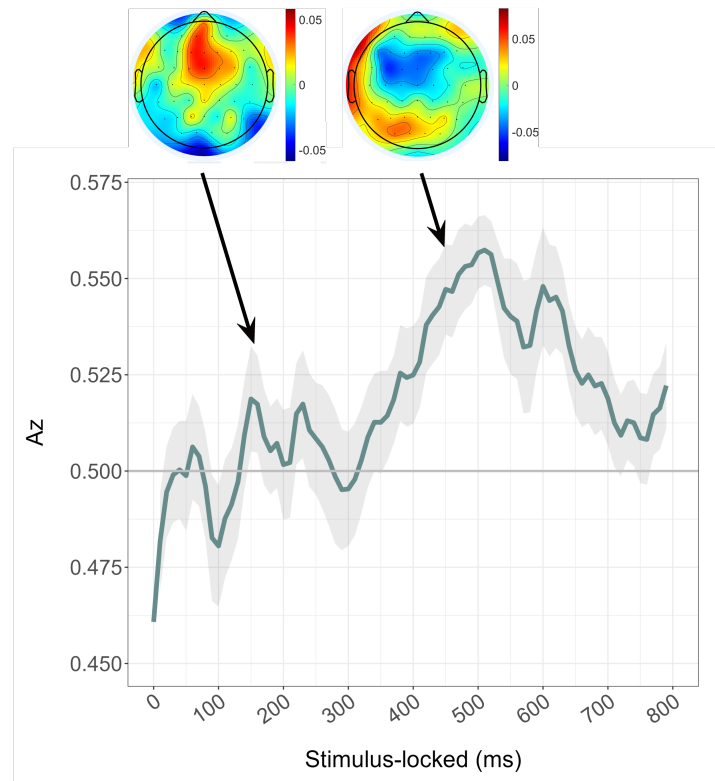


Figure 3.13 Discriminator performance. Negative and positive (signed) stimulus-PEs locked to stimulus onset and average scalp topographies (150-250ms and 450-550ms). (Az-values obtained with a leave-one-out cross-validation, and SE ribbon)

After obtaining the group-averaged Az values, we considered a broad window from 100 ms to 700 ms post-stimulus-onset, in which Az values per participant were above chance performance (0.5). We averaged the forward model estimates (a) across participants at 400 ms – 600 ms offset window to obtain one topography consistent with the average best discrimination. We aimed to find a timepoint per participant that was most consistent with the spatial pattern of stimulus-PEs. To do this we chose the time point at which the scalp projection was closest to the spatial representation (the group average a). Figure 3.14 below shows the average scalp topography obtained from the best time points. Figure 3.15 shows a sample of five participants' topographies from their best peaks.




Figure 3.14 Scalp topography from best peaks averaged across participants.




Figure 3.15 Scalp topographies from best peaks from a sample of participants.

Taking the best timepoint per participant (in the scale -100 ms to 800 ms post-stimulus), we extracted the EEG data at this time point from all trials. We then weighted the EEG data with the corresponding spatial filter (w) at the specific time post-stimulus and obtained the Y values (with EEG-native scale in microvolts). As such, for each participant, we obtained component amplitudes that most closely reflected the separation between negative and positive stimulus-PEs by taking into account the spatial generators and timepoint when the signal occurs.

Then, we averaged the weighted data across the three signed stimulus-PE levels, resulting in three points per participant. Figure 3.16 shows the EEG amplitudes obtained from the discriminator (Y values) at trials used for training the classifier (orange and blue), and testing (grey). Each point represents one participant's average Y value. The Y values followed a parametric effect consistent with the modelled stimulus-PE levels according to a one-way ANOVA analysis and Bonferroni corrected multiple comparisons between negative and middle values ($t(32) = 8.37$, $p < 0.001$) and middle and positive values ($t(32) = 8.5$, $p < 0.001$).




Figure 3.16 Y values (in micro-volts) plotted for each set of stimulus-PEs. Positive and negative PEs were used for training the classifier, while middle values were used to test the discriminator-resulting vector (w).

3.3.5 Predicting risk-taking from EEG amplitudes at signed stimulus-PE

Next, we obtained EEG amplitudes representative of the signed stimulus-PE according to the above results for each individual and each trial. We z-scored the Y values from all trials, ordered them in ascending order and binned them into 5. Figure 3.17 shows the probability of making a risky choice for each bin, with bin '1' being the lowest Y values, and bin '5' – the highest Y values. We fitted a logistic mixed effects model (estimated using ML and Nelder-Mead optimizer) to predict risky choices with Y values (formula: $p(\text{risky}) \sim Y$). The Y amplitudes (EEG stimulus-PE activation) at each trial significantly predicted the probability of choosing a risky versus a safe stimulus. The model included Y as random effects (formula: $\sim 1 + Y \mid \text{ID}$). The model's total explanatory power was weak (conditional $R^2 = 0.03$), and the part related to the fixed effects alone (marginal R^2) was 0.00564. Within this model: The effect of Y was statistically significant and positive (beta = 0.11, 95% CI [0.04, 0.18], $p = 0.004$; Std. beta = 0.14, 95% CI [0.05, 0.23]).

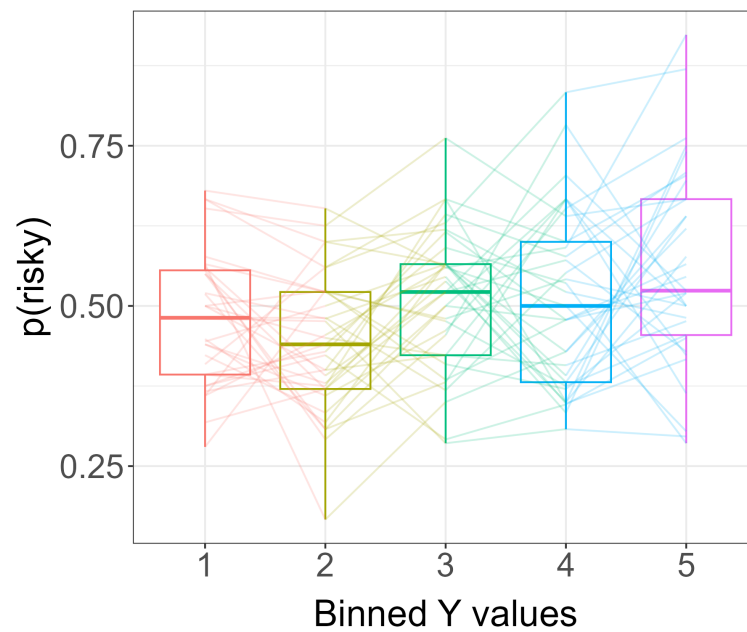


Figure 3.17 Binned Y values in ascending order with boxplots of the probability of risky choices. Each participant's points are linked with lines.

3.3.6 Exploratory analysis

As prediction errors are a key component of reinforcement learning in the brain, one may expect to find different generators between the different types of PEs (signed and unsigned), however, as this signal becomes more relevant for taking action, and encoding the environment, subsequent scalp projections may converge to similar patterns measured by EEG. Although not the main objective of this work, we explored the outcome prediction errors of the current experiment to test if they would be comparable to the stimulus-PEs in terms of spatial projections, something that has not been tested before. And additionally, we aimed to establish their relevance for the vast literature on feedback reward prediction errors within human studies.

3.3.7 Outcome-PEs

For the outcome-PEs we tested both signed and absolute outcome-PEs. We classified negative versus positive signed PEs and separately low versus high absolute PEs. Figure 3.18 below shows the results from these two classifiers with their most distinct topographic maps. For absolute PEs there were two clear topographies which we show as average between 250 ms and 450 ms, and a second one between 550 and 750 ms. For signed PEs, the earlier topography was more pronounced and sustained across time, the average of which was obtained from 100 to 300 ms, while the second one averaged over 400 to 660 ms.

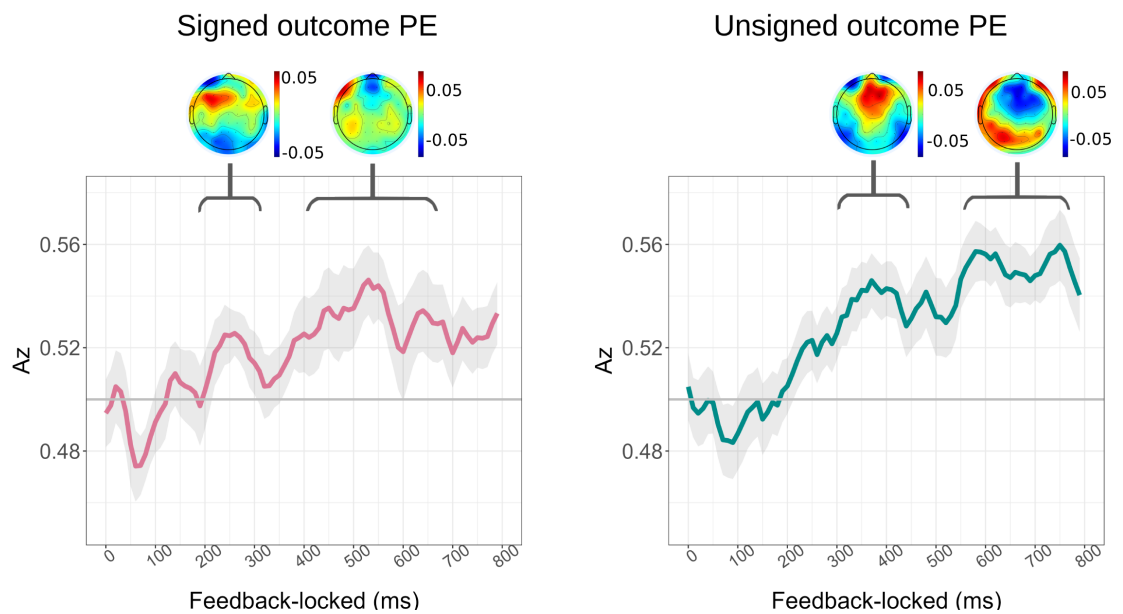


Figure 3.18 Az (leave-one-out) discriminator performance for Signed and Absolute outcome-PEs and topographic maps averaged across peak times.

3.3.8 Relating stimulus-PEs to outcome-PEs

Next, we compared the prediction error signals at the stimulus presentation and feedback presentation to look for similarities in the generators of the PE signal at both time points. We correlated the spatial projections (a) from the signed stimulus-PE with the absolute outcome-PE and the signed outcome-PE (each map containing 64 channels). Specifically, we extracted the spatial projections from the best peaks per participant and averaged the topographies which we then correlated (Figure 3.19), resulting in a positive significant correlation of $r(62) = 0.61$, (95% CI [0.83, 0.93]), $p < 0.001$, between stim-PE and absolute outcome-PE (Figure 3.19 left panel), and a small positive correlation of $r(62) = 0.3$, (95% CI [0.11, 0.55]), $p < 0.006$, between signed stim-PE and signed outcome-PE (Figure 3.19 right panel). As such, the signed stimulus PE was more similar to the unsigned outcome-PE than the signed outcome-PE in terms of spatial projections.

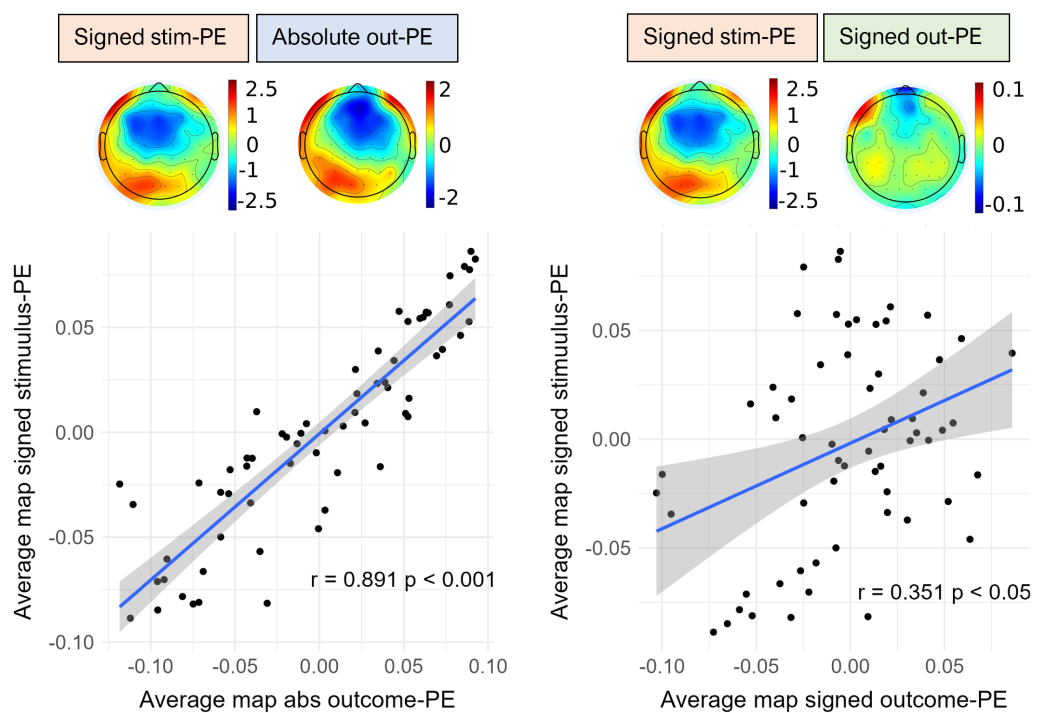


Figure 3.19 Correlations of topographic maps. Correlations shown between the spatial projections for signed stimulus-PE and absolute outcome-PE (left) and between signed stimulus-PE and signed outcome-PE (right).

3.4 Discussion

We tested a novel model based on learning in the basal ganglia, and specifically risk-taking following positive PEs. We replicated previous behavioural and modeling results and showed that the addition of a parameter that can account for individual differences can slightly improve the model fit. Our results supported the prediction that stimulus-PEs drive risky choices. Additionally, we showed an EEG component related to positive compared to negative model-based stimulus-PEs. This component appeared around fronto-central electrodes and showed higher amplitudes for positive compared to negative PEs at about 500 ms post-stimulus onset. It also significantly predicted risk-taking on a trial-by-trial basis, however, the effect was weak. Our exploratory results further showed the patterns of positive compared to negative (signed) outcome-PEs and low versus high (absolute) outcome-PEs. It was revealed that signed stimulus-PEs were generated in a fronto-central pattern similar to the absolute outcome-PE.

3.4.1 Cause of risk-taking

We found support for the hypothesised association between prediction errors and risk-seeking behaviour, and showed the neural pattern that is most closely related to this association. These results indicate that in line with a model based on the basal ganglia, dopamine functions may interfere with each other, given the computational modelling that we employed. Namely, based on the mechanistic model of risk-taking in the basal ganglia, dopamine as a teaching signal may affect the function of its risk-modulation properties. These results are a significant addition to the theory of risky behaviour and the model provides a possible cause that has not been considered until now.

It was previously found that reward-associated cues drive dopamine activity according to the estimates of future rewards, defining dopamine as a ‘teaching signal’ (Lak et al., 2014; Schultz, 2017a). Action demands, such as the need for a focused reward-seeking state, can shape Dopamine activity and subsequently increases in dopamine can invigorate responding (Grima et al., 2022; Robbins and Everitt, 2007; Wassum et al., 2013). The basal-ganglia risk-taking model suggests that it is the interplay between the two functions of dopamine that modulates risk-taking (Moeller et al., 2021). Our findings were best explained by the PEIRS2 model, which was built on the basal ganglia model and was able to explain the differential risk-taking behaviours depending on the preceding stimulus-PEs. It is possible that other models would show similar results in relation to predicting risk-taking via estimated stimulus-PEs, however, PEIRS2 is the only model that refers directly to a brain mechanism and allows further neurophysiological testing. Previous studies point to the influence of frontal cortices on the speed and accuracy of decisions (Lo and Wang, 2006) and it was shown that estimates of higher average rewards lead to increased vigour of motor actions (Guitart-Masip et al., 2011). This is consistent with our results showing that reaction times were faster after the presentation of the two stimuli with high EV (both-high) compared to low EV (both-low). We ruled out other theories of risk-seeking in our task, however, other models may be able to predict risk-seeking similarly to PERIS2. Next, I will outline important considerations for designing experiments aiming to test the PEIRS2 model and compare it to other theories.

Behavioural economics and finance have produced a few theories aiming to explain risk-seeking in humans. One such theory proposed humans are holding the mainstream theory of risk-reward trade-offs, taking risks when high rewards are sought, and avoiding risks at the expense of reward. As a well-known trading pattern, high risk (or variance) carries more chance of higher rewards, whereas low risk is more likely to lead to lower but safer outcomes. Behaviour related to the knowledge of risk-rewards trade-off was also related to a neural response in the insula during investments, which is a simple model and explains choices better than standard utility theory (Mohr et al.,

2010). The risk-reward trade-off theory predicts a person to be risk-avoiding (due to current goals or personality) when presented with a risky versus safer option. However, we did not find consistent risk-taking or risk aversion across all trials, suggesting that a risk-reward trade-off could not explain our results.

On top of the common knowledge of the risk-reward trade-off, humans also differ in their risk-taking due to long-lasting personality characteristics, conceptualised as risk propensity, risk disposition, or variance preference (Joseph and Zhang, 2021; Mishra and Lalumière, 2011). However, our general risk-taking and risk-propensity questionnaire results did not explain risk-taking in the bandit task. Previously, however, it has been suggested that risk-taking in laboratory tasks does not correlate with self-reported risk-taking (Guan et al., 2020) and does not generalise across different tasks (Helfersen et al., 2018). Some of these discrepancies between laboratory tasks and real-life decisions may come from what Daniel Kahneman explains as the effect of ‘playing for peanuts,’ in which people are found less prone to avoid risk with smaller wins and losses, compared to when playing with larger sums of money (Kahneman, 2012). Nevertheless, although the rewards provided in our task were small, we were able to show that risk-taking may be a consequence of a very short-term contextual bias of PEs. The result that a model with a free parameter that may explain individual differences was better than the one without that parameter, points out that individual differences may influence how and when risk is taken. We did not test individual differences beyond anxiety and general risk propensity, so we did not have the power to find what this parameter may mean, however, future studies aiming to incorporate individual and contextual differences should explore this further. This could be done by segregating participants based on the models that best fit their data, and collecting more individual data that could potentially explain why different models are better for some participants but not others.

For example, the s-shaped UTIL model was another model that fit the data considerably well across the span of the task and may be a better fit for some of our participants than the PEIRS models. The s-shaped UTIL is based on prospect theory, which predicts risk-averse choices for wins and risk-seeking within a loss context (Kahneman and Tversky, 2018). Within our task, losses can be considered as outcomes below the reference point of 50. We, however, found risk-seeking for potential gains (in the context of high EV cues) and risk-aversion for potential losses. This discrepancy between our findings and the predictions of prospect theory can be explained as a consequence of the description-experience gap (Moeller et al., 2021). Or, in other words, experiments that give a full description of available options lead to different behaviours compared to experiments in which cue-outcome associations are learned on the go (Wulff et al., 2018). As such, risk-seeking as explained by prospect theory tends to reverse if probabilities are learned on the go, and studies showed that risk-seeking increases in the context of wins if learning is incremental and based on feedback (ibid.). Consistently, there may be different neural mechanisms at play during risk judgments of known probabilities of rewards and those that have to be learned and may remain somewhat uncertain until the end of a task. For example, it was previously found that when uncertainty unravels over short time periods (and probabilities have to be learned on the go), the insula, the dorsal PFC and posterior parietal cortices are activated post-stimulus onset, while medial PFC is more active when probabilities are known beforehand (Huettel et al., 2005). Both known cues and cues learned on the go were implemented in one single study by Fitzgerald et al. (2010). In one condition cues explicitly described the probability of rewards and in the second – cues had to be associated with a probability distribution through learning. The ACC, vmPFC/mOFC were found active during choices of cues that require online risk learning, and the anterior insula and the inferior frontal gyrus/IOFC were active during risk judgments of known probabilities. Therefore, the involvement of the frontal cortex areas and PE biases in risk-taking may depend on the demands of the task. Contrary to our results, a study that used simpler cue-reward associations, with just binary options associated with each bandit, found risk aversion in the gain as well as loss contexts (Niv et al., 2012). The study also found that neural

activity in the vmPFC was sensitive to both the magnitude and probability of potential outcomes, with greater activity observed when potential gains were larger and when the probability of a positive outcome was higher. These findings are consistent with the previous studies that found vmPFC activation for learning risks on the go. In order to resolve whether the differences in risk-seeking are due to the complexity of cue-outcome associations, future research should incorporate cues of varying descriptive power including cues that predict more complex distributions of rewards and those that predict binary options, as well as explicit cue-probability mappings.

3.4.2 Spatio-temporal brain patterns of stimulus-PEs

Our results showed two pronounced components expressed at different times and with different topographies. We decided to focus on the second one as it was more consistent across participants, leading to higher average classifier performance. However, it is worth discussing the significance of both components in relation to previous ERP and other EEG studies.

ERP studies have identified a few signals that have been found in the time window between the presentation of a cue (/ warning stimulus) and a response. These include the cue-N2, cue-P3, contingent negative variation (CNV), and stimulus-preceding negativity (SPN). The cue-N2 has a frontocentral topography and is observed between 250 and 330 ms post-cue. It has been related to cue evaluation, showing stronger negativity for punishing cues (and possible loss) compared to rewards (Novak and Foti, 2015). Cue-P3/P3-W is a response that appears between 300 ms and 600 ms post-cue and is expressed as a stronger ERP for cues predicting higher compared to smaller rewards with a centro-parietal positivity (Broyd et al., 2012; Goldstein et al., 2006). The cue-P3 has been mostly related to the salience of the stimuli (Novak and Foti, 2015). It is thought to reflect the brain's response to the allocation of attention, and motivational processes, and was related to the ventral striatum in a simultaneous EEG and fMRI study (Pfabigan et al., 2014). The CNV, a slow negative potential over

fronto-central sites was associated with two waves, one related to orienting, and another to expectancy, while the stimulus-preceding negativity (SPN) appears during the feedback-anticipation period (Zhang et al., 2017). The most relevant for our stimulus-PE ERPs are the cue-N2 and cue-P3, as these ERPs are related to the evaluation of the cue rather than the preparation of a motor response or an expectancy of a reward. And while the cue-P3 has been found mostly insensitive to cue valence, the cue-N2 depends on the sign of the PE. Our second component is more consistent with the cue-N2 not only due to both capturing the sign of the PE but also in terms of the topography of fronto-central activation. This fronto-central activity may be related to areas relevant to the coding of the stimulus-PE. Namely, fronto-central EEG activation can be the result of medial PFC, ACC, OFC, or other adjacent areas (Hauser et al., 2015). These sources, which have to be confirmed by fMRI studies, would be consistent with the proposition that the stimulus-PE is a ‘bias’ from prefrontal areas on the ‘Pavlovian’ function of the basal ganglia. The vmPFC, OFC, and ACC receive projections from dopaminergic areas, and the PFC feeds back to the striatum (Chaua et al., 2018; Volkow et al., 2017). Although, due to differences in analysis, our methods cannot be completely compared to ERP studies, previous results of classification have found similar topographies as a response to valanced cues.

One such study that used valence cues and has a design that allows a comparison with EEG topographies related to reward PEs used simultaneous EEG-fMRI and tested how participants make choices between cues that probabilistically predicted binary positive or negative feedback (tick or a cross) (Fouragnan et al., 2015). The response to the feedback was expressed as two components, the first one was associated with an orienting response, or a type of surprise, which was stronger for cues predictive of negative compared to positive rewards (unsigned PE, Fouragnan et al., 2017), and hypothesised to be underpinned by noradrenaline in relation to studies showing a tight link between noradrenaline and surprise (Preuschoff et al., 2011). The second component predicted value updating on a trial-by-trial basis and was hypothesised to be associated with rewards and dopamine and was later interpreted as a signed-PE signal (Fouragnan

et al., 2017). As the feedback of the task by Fouragnan et al. (2015) was binary (ticks or crosses), our findings can be compared at cue onset as the pairs of cues we used for training the classifier were either more than an expected average or less. We found two components with similar topographies to the ones found by Fouragnan et al. (ibid.) but with different timings. In the original study, the components were found between 220 ms and 320 ms post-feedback, while we found one distinct topography around 150 ms and 250 ms and another one between 450 ms and 550 ms. The difference between the two studies is that the stimuli of our task were more complex in that they predicted a full distribution of rewards rather than a binary outcome, which may have led to more time needed for the appraisal of the sign of the PE.

3.4.3 Prediction errors post-stimulus and post-outcome

So far two models have been used to associate prediction errors with learning. One option is for learning to occur due to a violation of expectations, while another option is for learning to happen as a response to confirmation of expectations. The Pearce-Hall model posits that PEs drive learning through surprise. It predicts more attention and learning during events that result in overall larger (unsigned) PEs, hence, violations of expectations. Under this model, the magnitude of the reward is the key component that drives learning. Studies have shown that noradrenaline and dopamine released from the locus coeruleus stimulate hippocampal memories and place cells reorganisation in response to novelty that can be positive or negative in the form of food or shocks (Kaufman et al., 2020; Wagatsuma et al., 2017). On the other hand, dopamine-assisted transfer from a stimulus to a cue leads to increased dopamine firing when a reward is higher than the cue-predicted level, and a decrease in dopamine when the reward is lower than the learned predictions (Schultz, 2017c). This mechanism is in line with the Mackintosh model in which when a prediction error occurs, attention is shifted towards the cues that are most relevant for predicting rewarding outcomes. The Mackintosh model predicts enhanced attention and learning with cues that actually predict rewards

well (rather than negative outcomes). This is in line with studies showing the decay of learning under punishment compared to rewarding feedback (Nakatani et al., 2009) and memory consolidation in the absence of motivation or emotional arousal during memorising (Murayama and Kitagami, 2014).

We showed that the patterns of brain responses locked to the presentation of the feedback (outcome) were comparable to those related to the cue. Interestingly, the signed stimulus/cue-PEs were more like the absolute outcome-PE than the signed outcome-PE in terms of spatial projections. It is possible that part of the signals underlying the two components is related to learning. In terms of memory, both signed and unsigned PEs may be associated with increased dopamine to the hippocampus driving episodic memories, although in the case of unsigned PEs this can be additionally modulated by noradrenaline from the locus coeruleus (Rouhani and Niv, 2021). Consistent with our results that signed cue-PEs and unsigned outcome-PEs have similar spatial projections in the EEG signal, it was recently found that these PEs and their specific occurrence in time within a trial contribute similarly to learning. Rouhani and Niv (*ibid.*) found that on one hand, the unsigned PE at outcome modulated trial-specific learning rates, in line with a Pearce-Hall model focusing on the surprise of the PEs and, on the other hand, the signed PE at cue-onset predicted learning in line with a Mackintosh model. Future studies should establish if indeed the two types of PEs are associated with learning and if our findings are the result of this association. Although these PEs likely arise via separate areas including the basal ganglia, amygdala, locus coeruleus, and prefrontal areas (Murty et al., 2012), it is possible that the scalp projections we observed are the result of the convergence of different signals to similar cortical areas that are relevant for cue/feedback evaluation and motor action. Future simultaneous EEG and fMRI studies could look at the generators of the different prediction errors as the decision unfolds to reveal a more precise image of whether these generators are similar or different at cue and feedback evaluation.

3.4.4 Limitatons and future directions

Although our study is based on a mechanistic model of differential dopamine functions, we cannot confirm whether dopamine or its influence (or other neurotransmitters) is what drives the observed results. Moreover, the exact mechanism of dopamine on the two pathways is still under debate, especially due to the complexity of neural connections and differential dopamine affinity of D1 and D2 receptors (Kirschner et al., 2020). Although we did not measure dopamine levels, the topography we found could be interpreted as a result of medium frontal cortex projections that are highly related to dopamine function during reward-related tasks. These findings, however, should be interpreted with caution due to the small effects, and future studies aiming to test the theory mechanistically should either modulate dopamine levels within the two pathways as done previously (Jenni et al., 2017) or should measure dopamine within single trials.

Although our task design is higher in ecological validity, as it utilises more complex cues in the form of continuous distributions, a limitation of employing such novel paradigms is that they are less comparable to previous studies that have used cues indicating binary outcomes. For example, it is to be further studied if the predictability of cues (predicting simpler or rewards with higher entropy) results in different temporal profiles of these reward-specific components. Additionally, our classification analysis is powerful in that it uses the full set of EEG electrodes and a high amount of samples across trials. This, however, may hinder the comparability between ERP studies and those using more complex methods. Although the harnessing of computational methods and machine learning techniques is desirable, future studies focusing on ERPs should determine if indeed the cue-P3 is more associated with signed or unsigned PEs. This may help answer questions about the parallels between findings from LDA and the bulk of existing ERP studies.

Our study is an important addition to the value-based decision-making research, showing that risky decisions can be influenced by short-term fluctuations of expected rewards and that these expected value fluctuations can be tracked in the patterns revealed by EEG. Although we cannot infer that trial-by-trial increased levels of phasic dopamine are the cause of risk-taking, we showed that this is the most likely explanation given the models we tested. As such, future research should test the connection between the influence of positive and negative reward prediction errors on choice by measuring dopamine and cortico-basal ganglia connectivity.

Group Membership Bias and Stimulus Prediction Errors Influence on Risk-Seeking

In the previous chapter, we showed that a model based on learning in the basal ganglia can predict risk-taking on a trial-by-trial basis and that this behaviour can be explained by the neural response to stimulus-PEs. In this chapter, we build on these findings and add an additional potential bias that was proposed to come from the mere observation by in-group or out-group members. Based on social identity theory we expected that risk-taking would increase in the in-group condition followed by positive reward PEs. We found that differences in overall risk-taking between the two social conditions could not be distinguished, however, they differed from the baseline. Additional measures of mental health showed that the results can partially be explained by different anxiety levels. We discuss possible explanations of the results ranging from exploration to increased arousal and potentially elevated stress levels due to the observation procedure.

4.1 Background

Social identity is the concept that explains differential behaviours and attitudes when we identify with one group or another (Myers, 2016). Group membership can change self-identity, feelings of confidence, and the need to be approved or liked by members of the ‘in-group,’ defined as the group that we feel close to and part of (Voci, 2006). Subsequently, different behaviours (than when alone) can stem from perceived norms about what the individual thinks is appropriate for the group (Abrams et al., 1990), or simply from feelings of diminished responsibility and self-awareness (Prentice-Dunn and Rogers, 1982). Several studies have found that individuals tend to take higher risks when interacting with members of their in-group compared to out-group members. This phenomenon is known as the “in-group advantage” or “in-group favouritism” and has been observed in various contexts, such as financial decision-making, negotiation, and cooperation.

For example, participants were more likely to take financial risks and make risky investment decisions when interacting with members of their in-group, compared to out-group members (Hertel et al., 2000; Hogg et al., 1995). This suggests that individuals may feel a higher level of trust and cooperation with their in-group, which can lead to reduced uncertainty and increased risk-taking behaviour, supporting the social identity theory model of risk-taking (Cruwys et al., 2020b, 2021, 2020a). The effects of group membership in the studies by Cruwys et al. (2020b, 2021) were seen even with arbitrary criteria used for groupings such as using the scores on colour perception or dot-count estimation tasks. And the effects were amplified by *integrity beliefs*, the degree to which others were perceived as dependable and honest.

However, there is also evidence that due to inter-group conflict individuals may take risks in out-group conditions, particularly when the risks involve competition or the need to establish dominance, or under threat of resource depletion (Safarzynska, 2018). Additionally, other factors such as individual differences and mental health conditions

can also influence risk-taking behaviour in in-group and out-group conditions. Studies by Dreu et al. (2010) have shown that individual differences in personality traits, such as extraversion and agreeableness, can also influence risk-taking behaviour in group settings, while mental health can affect the representation of prediction errors, risk-taking, and exploration.

Anxiety is a mental health condition which affects daily living and functioning and has a prevalence of about 30% in the USA, and about 7% in Europe (NIH, 2022). Anxiety is mostly expressed as a heightened response to uncertainty, which involves a psychophysiological cascade of negative events. It is possible that the response to unexpected events is inflated, leading to a higher physiological or emotional response than to less-surprising events. The brain area that converges different domains of prediction errors may be the insula. Interoceptive alterations (own bodily perceptions) were found to be underpinned by the insula (Stein et al., 2007). However, apart from monitoring bodily processes, the insula was also found to follow predicted risk and prediction errors (Preuschoff et al., 2008). One may speculate that processing the variance of rewards and the variation in bodily states may be two aspects of a general process tracking the distribution of prediction errors.

Several computational frameworks attempt to explain how anxiety affects decision-making and learning. A recent study modelled the trade-off between stochasticity (small and more frequent) variations of rewards and volatility (large and unexpected changes) of rewards (Piray and Daw, 2021). According to reinforcement learning larger prediction errors, as in the case of volatility, lead to a higher necessity to update currently held states, hence, they drive learning faster than small prediction errors. The study by Piray and Daw (*ibid.*) study proposed that estimates of prediction errors may be altered by anxiety, and this may also in turn change learning and decision-making. The study showed that higher anxiety can predict faster learning due to an over-amplification of prediction errors. Additionally, although anxiety generally reduces risk-taking and

exploration (Giorgetta et al., 2012), increased conformity in anxious individuals can sometimes lead to them undertaking undesirable behaviours when in a group (Lewis et al., 2008). With the current study, we aimed to explore whether anxiety would increase or decrease risk-taking modulated by social factors.

Despite the advances in the science of decision-making, it is still unclear how social factors and different levels of social interaction (from mere presence to competition/cooperation) influence decision-making under uncertainty, and how this is translated into neural signals. It was previously found that even a simple grouping manipulation is enough to induce quick feelings of being part of a group, which additionally led to the activation of different brain areas to in-group compared to out-group members (Hobson and Inzlicht, 2016; Bavel et al., 2008). Brain areas involved in non-social decision-making have also been implicated in social-decision making, suggesting a common currency schema as opposed to a social-valuation-specific schema (Ruff and Fehr, 2014). The basal ganglia (including the striatum), different parts of the prefrontal cortices, the ACC, and the insula have been implicated in value-based and risky decision-making. All of these structures have also been found in similar processes of decision-making within a social context. The anterior insula and the ACC, previously implicated in salience coding, were found more active during decisions to financially benefit others compared to harm them (Greening et al., 2014). Also, in a study where participants watched baseball videos, the ACC and anterior insula were more active after own-team failure compared to loss by the opposite team (Cikara et al., 2011). Opposite to that, the same study found higher activation in the ventral striatum when the own team succeeded, and the other team failed. This was supported by another study showing higher ventral striatum activation for out-group members receiving misfortune or pain than in-group members (Takahashi et al., 2009). And the OFC was more active when participants in a study saw novel in-group members compared to novel out-group members (Bavel et al., 2008). The above studies show

that there is a big overlap between the social and non-social brain areas when it comes to decision-making and vicarious experiences with different group membership contexts. However, it is still unknown if the brain's temporal dynamics during social and non-social decision-making are similar under uncertainty.

From a separate line of electrophysiological research, there is evidence that mere observation by an out-group compared to an in-group member can result in differential EEG patterns during decision-making. The mere presence of others can influence arousal due to fear of negative evaluation (Mullen et al., 1997) and can alter behaviour due to perceptions of others' beliefs and values (Stone and Allgaier, 2008; Suzuki et al., 2016). A research group at the University of Toronto tested participants being observed by someone either 'on their team' or not, with EEG (Hobson and Inzlicht, 2016). The participants in the study were made to believe they scored similarly to one of the confederates on a dot estimation task (in-group member) and had a very different score from the other confederate (out-group member). The feedback-related negativity (FRN) is a centroparietal ERP that peaks over centroparietal electrodes and around 300 ms to 600 ms as a response to feedback. It has been associated with prediction errors and activity in the ACC. In the study, the participants' FRN was attenuated while being observed by an out-group compared to an in-group member. These results suggest that even a simple grouping method can result in different brain responses during value-based decisions.

The above examples show the importance of including social influence in the model of risk-taking to increase the ecological validity of laboratory testing and explore the generalisability of findings across domains. In the current study, we aimed to combine the influence of group membership, anxiety levels, and trial-by-trial prediction errors to test their effect on risk-seeking. We hypothesised that risk-seeking will be higher

in the in-group condition compared to the out-group condition and that this would manifest as an additional bias on the already present bias of positive stimulus-PEs on risk-seeking. In relation to anxiety, we did not have concrete predictions and treated the analysis as exploratory.

4.2 Methods

The materials and task are extensively explained in Chapter 3 Methods section. Within this chapter, we focus on the social manipulation that was added to the experiment to test the influence of the mere presence of an in-group and an out-group member on risk-taking behaviour and the associated neural responses.

4.2.1 Social conditions

To test in-group and out-group membership bias on risk-taking, we added two social conditions following a baseline block. Participants completed 180 trials of the task on their own, and in two blocks of 180-trials each, they completed the task while believing to be observed by either of two other participants, one from their in-group, and another from their out-group. To assign participants to their out- and in-group members, we asked them to estimate the number of dots in the image below (Figure 4.1).



Figure 4.1 Image shown to the participants to assign them into social groups. The participants were asked to estimate the number of dots on this image and were told that they were either over- or under-estimators. It was explained to them that during the task they would be observed by two other people who scored similarly to them (e.g. over-estimator) or with the opposite pattern (under-estimator).

Then, they were told if they were an under-estimator or an over-estimator. Like in the study by Hobson and Inzlicht (2016), our participants believed that one of the confederates had a perceptual performance like theirs (either an under- or overs-estimator), while the other confederate had the opposite result. The experimenter also assured the participants that the estimation they made on this task was not in any way related to another cognitive or perceptual ability and that all people either overestimate or underestimate the number of dots. This was done to ensure the participants do not experience doubt or a ‘stereotype threat’ (Steele, 1997).

The method is an arbitrary grouping strategy that was found to be enough to induce feelings of shared group membership (Hobson and Inzlicht, 2016). The participants were not aware of the actual reason for social grouping; however, they were told that the online observation aimed to determine how people learn by watching someone else perform the task, and how this was related to observing someone from their perceptual estimation category or the opposite one.

Participants were made to believe that two other people were observing their screen while performing the task. To achieve this belief, we created two videos of a person connecting via Zoom. Before the experiment began, the participants were seated in the EEG booth and observed the monitor being controlled externally by the experimenter who was preparing the software for the task. This part was essential to show the participants that they could see when the experimenter was operating the computer.

The pre-recorded video started with a Zoom screen appearing with the option to share a specific window on the computer. (An example of this procedure is shown in Figure 4.2). Then it appeared as if the experimenter moved the mouse and selected the screen of the task to be shared with the person on the other side of the call. The typical controls that are shown when a screen is shared in Zoom then slid from the top and a small window with the video of a confederate was shown on the screen. At this point, the task screen displayed the sentence: “Dear participants, Thumbs up, if you can see the screen.” and the confederate lifted a thumb in agreement (which was timed very precisely to appear as a natural reaction). After the confirmation, it appeared that the experimenter dragged the small window showing the confederate outside of the task screen as if transferred to an adjacent monitor. Then it looked like the Zoom screen-share controls were hidden by the experimenter clicking on a small ‘hide controls’ icon. An indication that the screen was still shared was a green frame spanning the edges of the full screen and lasting throughout the whole block.




Figure 4.2 An example screen of a pre-recorded Zoom connection with a confederate. The participants observed someone being connected via Zoom and showing thumbs up, indicating that the screen-sharing is working. The participants were made aware that the computer in the lab does not have a camera, so the other person would not be able to see them.

4.2.2 EEG analysis

As explained in the previous chapter, the model-based stimulus-PEs that we obtained from fitting the best model, were divided into three equal parts of 60 trials each per block. The linear discriminator was then trained with the EEG data associated with the negative and positive stimulus-PEs, while we tested the resulting vectors with the testing data that were the middle trials. Similarly, for the outcome-PE, we divided trials into negative and positive outcome-PEs, obtained from the model fits, and trained a separate classifier on these trials. For any LDA results, we considered anything above 50% (chance level, or A_z of 0.5), as indicative of significant discrimination. We obtained the spatial projections of each component via a forward model.

4.2.3 Behavioural analysis

As detailed in Chapter 3, the behavioural data for the social conditions were analysed in the same way as the baseline. The data was divided into trials with both-high, both-low, and different means. As we were interested mainly in risk-taking between stimuli with the same means, only these trials were used for some of the analysis. We performed t-tests on the social blocks to test the proportion of risky choices in both-high compared to both-low conditions. We also divided each block into 5 bins (36 trials for each bin) and computed the differences between both-low and both-high risk-taking on each bin. Like in the baseline block, reaction times were positively skewed, and two mixed-effects models were used, one for each social condition. Participant ID was input as a random factor, pair type (both-high/both-low/different) as a predictor, and reaction times in milliseconds as an outcome. This was done using the *afex* package in R (Singmann et al., 2016).

4.3 Results

None of the participants was able to tell that the videos were fake and that the people who were connected via the pretend Zoom call were confederates. This was revealed by a questionnaire administered at the end of the experiment prior to debriefing which asked the participants to rate the quality of the Zoom connection, and to report the ‘feeling towards the same-group member and opposite group member.’ Verbally when asked after the experiment all participants said they thought the Zoom connection was good. The question about the Zoom connection was an implicit prompt aiming to test if anyone would only then express doubts about the procedure, however, all of the participants focused on trying to recall if the connection was good without other comments. We took this as an indication that the procedure was successful in concealing the recordings as live videos. Each participant gave a rating for the in-group and out-group observers as ‘part of my team’. As these ratings were on an ordinal scale we tested the difference between ratings given to the in-group and out-group observers by each participant with a Wilcoxon signed rank test in R, showing higher ratings for the in-group observers, $V = 106$, $p = 0.008$, pseudomedian = 1.99 (95% CI [0.5, 3]). The average response to the question ‘Do you think the observer paid attention to you playing the game?’ participants answered on average with 6.67 (out of 10, $sd = 3$) for the in-group condition and 6.2 ($sd = 2.98$) for the out-group condition (with no significant difference between the two). A paired Wilcoxon signed rank test showed no significant difference between the two social conditions: data: $V = 27.5$, p -value = 0.587, pseudomedian = 1 (95% CI [-1.9, 2.9]). Figure 4.3 shows boxplots of the ratings. These results indicate great variation between participants, with some expressing very low belief that the other person observed their performance, and others indicating the opposite. As there was consistency within participants’ responses on the two social conditions, we could confirm the deception worked as intended and next, we turned to examine the behavioural and brain patterns between the two social blocks.




Figure 4.3 Ratings on belief that the in-group and out-group members observed the experiment.

4.3.1 Behavioural results

Our hypothesis was that participants would be more risk-seeking in the in-group condition, especially in the both-high set of trials, while less risk-seeking in the out-group condition. The scatter plots below would be consistent with the theory if more points were in the low right triangle of the plots for the in-group condition. Unlike the baseline condition, throughout both the in-group and out-group conditions, participants were not more likely to select a riskier option when presented with both-high EV stimuli, compared to both-low EV stimuli as shown in Figure 4.4. With in-group paired t-test resulting in: $t(32) = 1.98$, $p = 0.067$, mean diff. = 0.11 [-0.008 0.23]; and out-group $t(32) = 0.88$, $p = 0.38$, mean diff. = 0.08 [-0.09 0.22]. For the in-group condition, 58% of participants made more risky choices in the both-high condition relative to both-low, whereas this was true for 54% of participants in while in the out-group condition. A paired samples t-test showed no significant difference between the risk-taking at both-high condition for in-group compared to out-group context ($t(32) = 0.9$, $p = 0.3$, diff = 0.04 [-0.05 0.12]).




Figure 4.4 Probability of risk-taking in both-high and both-low trials across the social blocks.

Figure 4.5 reveals the average pattern of risk-taking across the span of the full task for each social condition, each bin consists of 36 trials. To determine if there were significant differences between risk-taking for both-high and both-low conditions, we conducted two separate repeated-measures ANOVAs for each social condition with participant ID as an error term, average risk-taking as the dependent variable, and pair type and bin (from 1 to 5) as factor independent variables. For the in-group condition, the repeated-measures ANOVA (formula: $\text{risky} \sim \text{factor}(\text{bins}) * \text{pair-type} + \text{Error}(\text{ID})$) suggested that: For the in-group condition the interaction between bins and pair-type was significant and small ($F(4, 120) = 2.61, p = 0.04; \eta_p^2 = 0.03$). For the out-group condition the interaction between bins and pair-type was significant and small ($F(4, 120) = 3, p = 0.02; \eta_p^2 = 0.03$). Additionally, Tukey's HSD was done to determine the Bonferroni adjustment of the p-value for multiple comparisons. This resulted in the significant differences between pairs shown in the plot as stars. There was a fluctuation between the risky and safe choices, however, for the in-group condition there was a higher risk-taking tendency during the very beginning of the task (36 trials) and after the middle of the block (last 72 trials). For the out-group condition, risk-taking was below 50% for both-high as well as both-low conditions until the middle

of the block (around the 100th trial), after which risk-taking increased for the both-high EV condition compared to the both-low. On the other hand, there was generally less risk-taking on both-low EV trials for both social conditions, being less than 50% throughout the blocks.

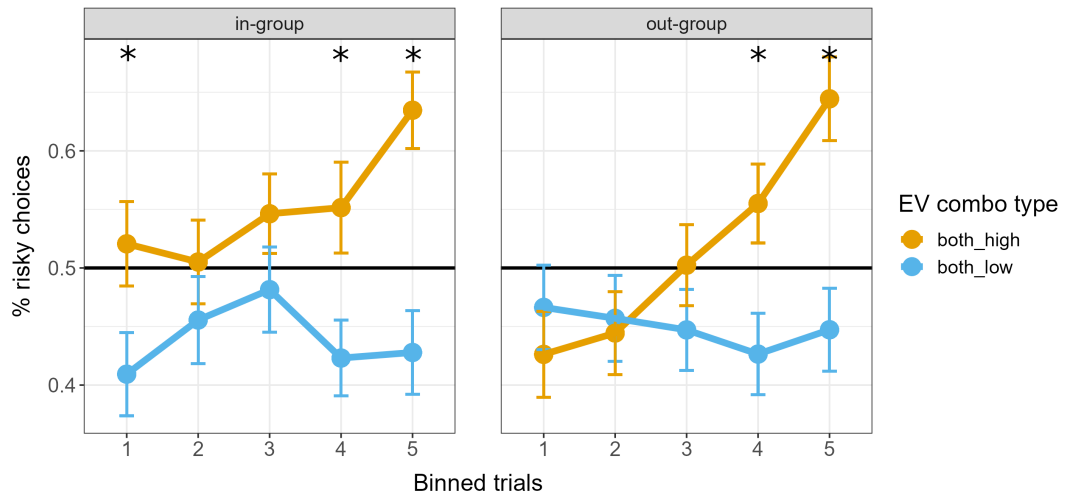


Figure 4.5 Learning curves for the social conditions. Choices of risky trials across the two blocks of in-group and out-group conditions. Each bin contains 36 trials (a total of 180 trials per block). Percentage of risky choices are shown as average with dots, with SE bars. The significance between the pairs, shown with stars, is indicative of $p < 0.05$, tested with a Tukey's HSD and adjusted for multiple comparisons.

Figure 4.6 shows the reaction times across blocks and conditions. Consistently with the baseline, for in-group and out-group conditions the reaction times were faster when a high-mean option was present as in the different-means condition, and when both stimuli were with high EV. The ANOVA main effect of pair type on RTs showed $F(2, 288) = 40, p < 0.001$. This result was independent of social condition with the interaction between pair type and social condition on RTs: $F(2, 288) = 0.13, p > 0.05$.




Figure 4.6 Reaction times are shown as average values per pair of expected values for each social condition.

We tested whether the performance differed between blocks. Performance was defined as the percentage of high-EV choices relative to low-EV in the condition where two stimuli of different means were shown on the screen. There was no significant difference between the performance of participants across conditions as shown in Figure 4.7 and tested with a repeated-measures ANOVA ($F(2,64) = 1.7, p > 0.05$). Due to the different risk-taking behaviours at the first and second half of the blocks, we additionally tested participant's performance after learning (30 trials) and dividing the rest of each block in half: into 1st stage after learning (30 to 105 trials) and 2nd stage (105 to 180 trials). Using a repeated-measures ANOVA we found no significant difference in performance between the 1st and 2nd stage of the trials after learning ($F(1,96) = 1.46, p > 0.05$).




Figure 4.7 Average performance across blocks with standard error of the mean. Performance was defined as choosing a stimulus with a higher mean, given two stimuli of different means.

4.3.2 Model selection

Due to the behavioural variability in the two social conditions, we used the full experiment to fit the models of interest to the data. The left panel of Figure 4.8 shows the differences in risk preferences across the two conditions of interest (both-high and both-low). S-shaped UTIL, PEIRS, and PEIRS2 were the models that produced simulated data that were not significantly different from the experimental data.

The right panel of Figure 4.8 shows the BIC values for each model summed over all participants. A lower BIC value indicates a better fit. The model with the lowest BIC was PEIRS2 (BIC = 7598) followed by PEIRS (BIC = 7695).

We next used PEIRS2 to model stim-PE and to obtain discrimination trials for the LDA done on the EEG data. Modelled stimulus-PEs significantly predicted trial-by-trial risk-taking (independent of stimulus pair types) across the task, as indicated by a mixed-effects logistic regression. Two models were fit for each condition with the glmer function in R, family = 'binomial', and the formula: $\text{risky} \sim \text{PE} + (1 + \text{PE} | \text{ID})$, in which risky was 0 or 1 depending on the choice, and PE - prediction error. For the in-group we obtained: $\text{beta} = 0.07$, 95% CI [0.03, 0.11], $p < .001$ and out-group: $\text{beta} = 0.05$, 95% CI [0.001, 0.10], $p = 0.044$.




Figure 4.8 Model comparisons for the full task. Blue bars indicate the simulated datasets with parameters obtained from each model fit (with SE error bars). The grey bar shows the original experiment results. The models that produce the closest results to the data are indicated as not significantly different (n.s.) from the original risk preferences. This was obtained using separate t-tests.

4.3.3 Model-based EEG results

Figure 4.9 shows the discriminator performance averaged across trials for the in-group and out-group blocks. The discrimination between positive and negative stim-PEs was consistent between the two social conditions, peaking at about 350 ms after stimulus onset and resulting in similar spatial projections as shown on the topographic maps.

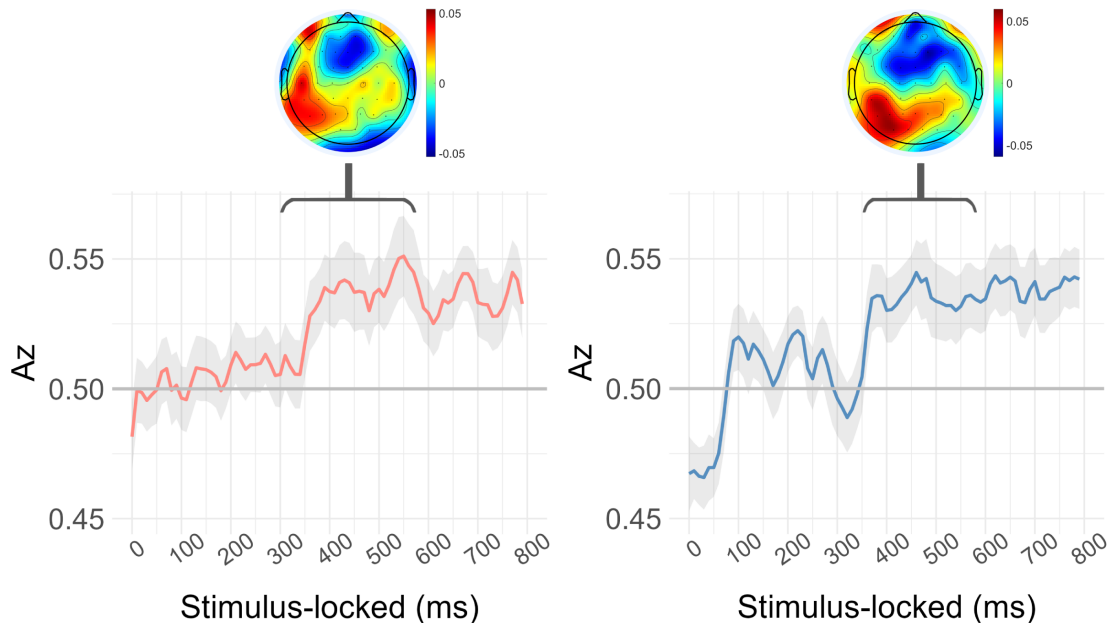


Figure 4.9 Discriminator performance and best maps. For in-group (left) and out-group (right) conditions. Comparisons made between negative and positive stim-PEs. The topographic maps were obtained from participants' best-points of discrimination (see Methods for details).

Next, we tested the discriminator by applying the weights from both-high and both-low trials on the EEG data on trials that we did not use to train the classifier. Figure 4.10 shows the result for the two blocks. The testing trials are indicated in grey and the training trials in orange and blue. There was a significant parametric effect in both conditions with amplitudes being lowest for negative PE trials, followed by middle trials, and highest for positive PEs. Repeated measures ANOVA results for in-group: $F(2,64) = 70.88$, $p < 0.001$, $\eta^2 = 0.62$; and out-group – $F(2,64) = 107$, $p < 0.05$, $\eta^2 = 0.77$.

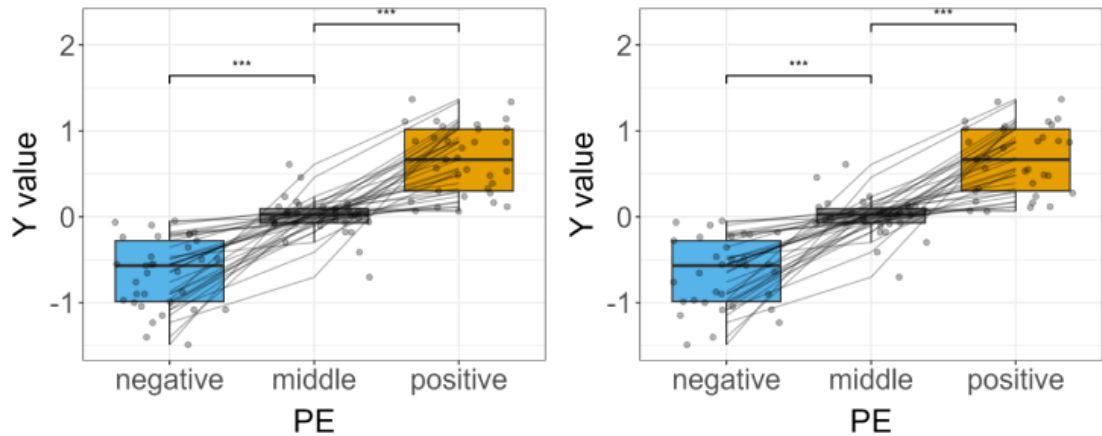


Figure 4.10 Average Y values for trained and tested (grey) trials. Left – in-group condition. Right – out-group condition.

Although the discrimination between negative and positive stim-PEs was encoded by participants (according to discrimination performance as shown in Figure 4.9), the trial-wise Y values did not predict trial-by-trial risk-taking. We tested this with a mixed-effects logistic regression using Y values and social condition as predictors, risk-taking (0/1) as an outcome, and a random intercept for subject ID across Y and social condition (formula: $\text{risk-taking} \sim Y * \text{social-condition} + (1+Y*\text{social-condition}|ID)$). The effect of Y on risk-taking was not significant (beta = 0.07, 95% CI [-0.05, 0.19], $p = 0.256$; Std. beta = 0.10, 95% CI [-0.07, 0.27]). And the interaction between Y value and social condition was not significant (beta = 0.07, 95% CI [-0.09, 0.23], $p = 0.371$; Std. beta = 0.10, 95% CI [-0.12, 0.32]).

4.3.4 Exploratory Anxiety effects

We measured state and trait anxiety with STAI (Charles et al., 1972). We aimed to test these in relation to a recent study showing that prediction errors may be processed and coded differentially for people with varying levels of anxiety (Piray and Daw, 2021), and hypothesising that this effect may be a mediator in risk-taking under social conditions. We found that anxiety scores as measured by SAI and TAI scales were significantly correlated with risk-taking during the ‘both-high’ condition for both in-group and out-group social blocks as shown in red circles in Figure 4.11. Interestingly, during the ‘both-high’ condition SAI and TAI scores were positively correlated with risk-taking in the in-group and out-group blocks, and during the ‘both-low’ condition in the out-group block anxiety was negatively correlated with risk-taking.

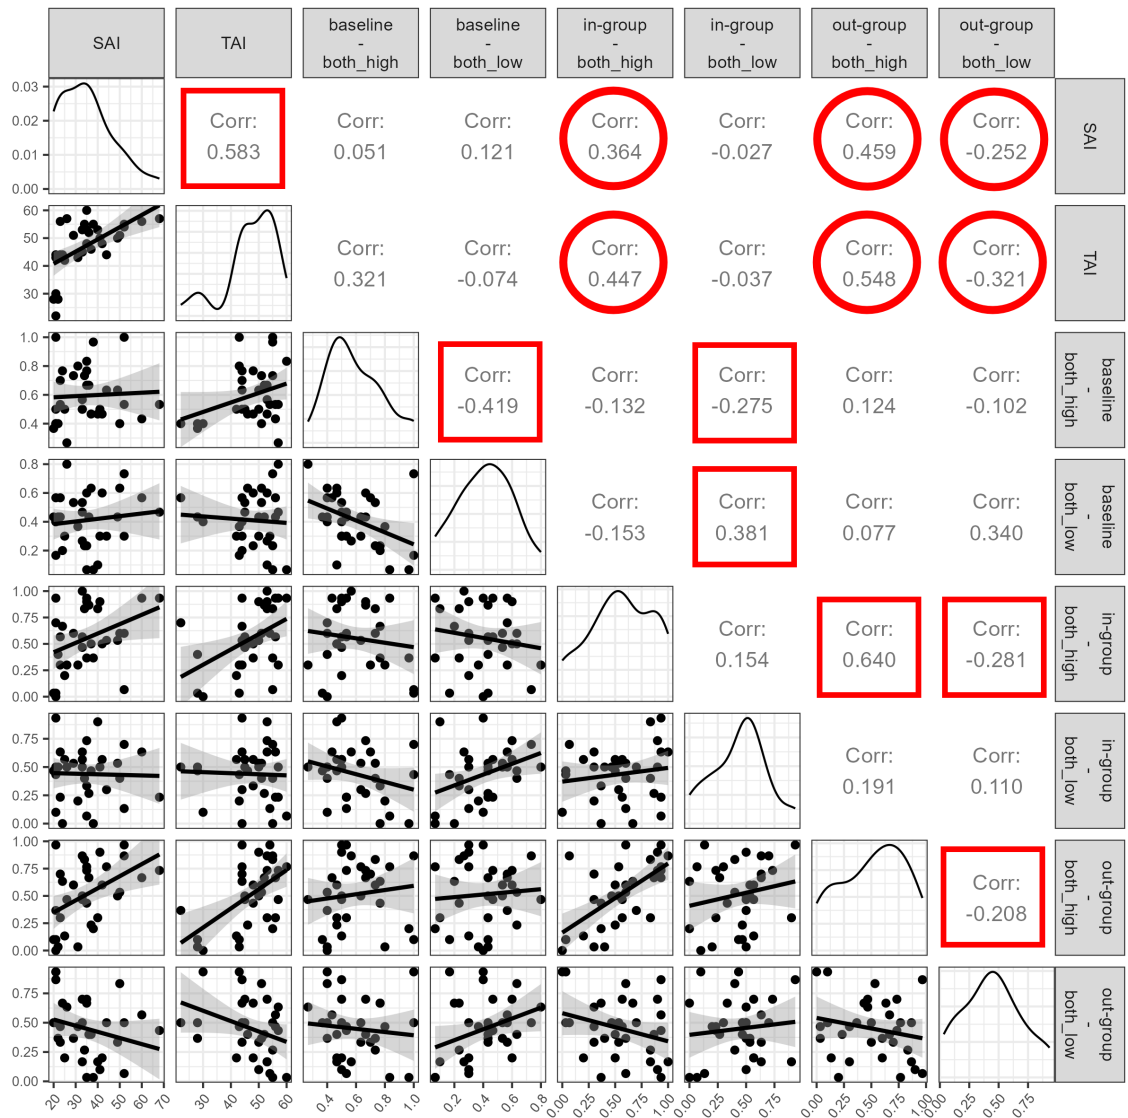


Figure 4.11 Correlation plot of risk-taking during baseline and social conditions and trait (TAI) and state (SAI) anxiety. With significant correlations shown with red frames. And significant correlations of interest are shown with red circles.

We focused on trait anxiety, as this was the focus in previous studies showing different prediction error processes between anxiety levels (discussed in detail in ‘Background’). We conducted a generalised linear mixed-effects model predicting risk-taking (0/1) on each trial with pair type and continuous trait anxiety score, with random slope and intercept for participant ID (formula: $\text{risk-taking} \sim \text{pair-type} * \text{TAI-score} + (1 + \text{pair-type} | \text{ID})$). We ran two such models on in-group and out-group social conditions separately. Overall trait anxiety predicted risk-taking in interaction with pair type for the in-group condition with a positive effect: (beta = 0.10, 95% CI [0.03, 0.17], $p = 0.004$; Std. beta = 0.93, 95% CI [0.30, 1.56]). Hence, there was more risk-taking during in-group condition with increasing anxiety for both-high means compared to low means. The same was true for the out-group condition with an effect of pair type * TAI-score on risk-taking being significant and positive (beta = 0.16, 95% CI [0.08, 0.23], $p < .001$; Std. beta = 1.41, 95% CI [0.73, 2.09]).

As was previously suggested, anxiety may improve performance on some value-based tasks. We also tested the performance expressed as more choices for higher mean options in the case of different means (Fig. 4.12). We found no effects of social block or anxiety, or the interaction of the two on percentage correct choices.




Figure 4.12 Percentage correct choices (of high EV within the different-means condition) across blocks.

As the differences in risk-taking that we found may be driven by different brain responses to stimulus-PEs, next we tested the relationship between anxiety and the Az or Y values across participants. In terms of neural activity, there were no significant differences between either the Az values or Y values across trials and participants with different levels of anxiety (average Y values shown in Figure 4.13, the main effect of anxiety was statistically not significant and very small ($F(1, 62) = 0.02, p = 0.897; \eta_p^2 = 2.70e-04, 95\% \text{ CI } [0.00, 1.00]$)).




Figure 4.13 Average Y activation across blocks, anxiety levels, and EV-pairs for signed stimulus-PEs. No significant difference between participants with different anxiety levels was found.

4.3.5 Response to wins and losses. Outcome-PEs in relation to social condition

We further aimed to replicate the findings of Hobson and Inzlicht (2016) who used an ERP analysis and found significant differences between win and loss trials mediated by the mere observation of in-group and out-group members. Wins and losses result in positive and negative PEs, driving either the Go or No-go pathways respectively. If a social bias acts on the basal ganglia and these PEs, it might translate as different activations between the two social conditions for the difference between positive and negative PEs (measured by our discriminator). Specifically, in the study by Hobson and Inzlicht (ibid.), there was a significant difference in ERPs between win and loss trials around 280 ms post-feedback for the in-group condition, but not for the out-group condition. As the presence of a difference is the typical signal, this result was interpreted as disrupted processing during the out-group condition. We ran two separate LDA analyses on the outcome-PE for the in-group and out-group conditions. We considered

any feedback above the grand average of 50 to be a win, and below that – a loss. Similar to the original ERP analysis, we focused the window of interest between 150 and 350 ms post-feedback. We looked for differences between the in-group and out-group conditions in the discrimination results as shown in Figure 4.14, however, the Az's were not significantly different (paired $t(32) = 0.03$, $p > 0.05$).

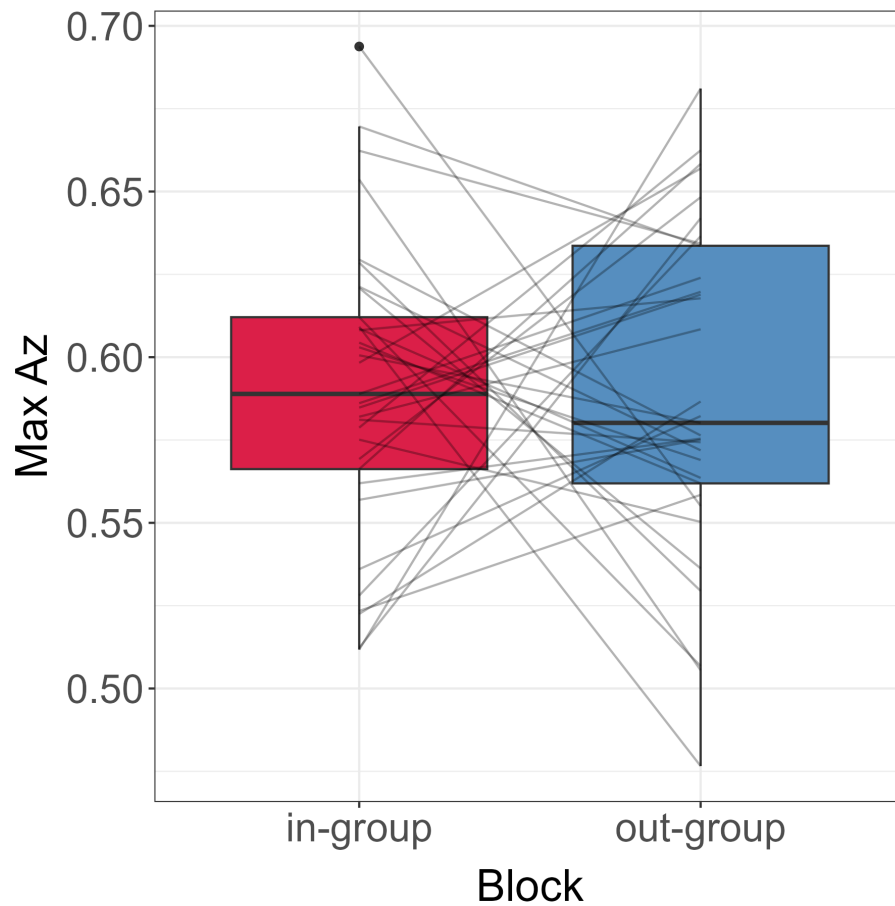


Figure 4.14 Boxplots of maximum Az values between 150 ms and 350 ms post-feedback onset for the two social blocks.

4.4 Discussion

In the current study, we tested the influence of the group membership bias on risk-taking as predicted by a model of prefrontal influence on the basal ganglia. We made the participants believe that an in-group and an out-group member observed them while performing the task to test the bias of mere presence on risky decision-making. We found that risky choices during the social conditions changed notably from baseline, showing an initial decrease in risk-taking, followed by increased risk-taking for both-high trials. The EEG response to stimulus-PEs in the social conditions was comparable to the baseline in terms of temporal and spatial profile, although the extracted component could not predict risk-taking during the social conditions. Additionally, we tested the influence of anxiety on risky choices within in-group and out-group conditions. Our results showed increased risk-taking in participants with higher anxiety for both social contexts and specifically during the condition with higher means. Combined, the above results are in line with a decision-making pattern that is likely the consequence of a combination of social context, stress, and model-free versus model-based mechanisms.

4.4.1 Social bias or exploration

Given the design of our experiment and the instructions given to the participants, the results we observed in risk-taking may be the consequence of a few conscious decisions being implemented or the influence of implicit biases. The participants may have decided to 1) keep playing the game as if no other person was ‘present’, 2) employ a different strategy that would benefit the in-group member if they tried to learn from their choices and hinder the performance of the out-group member, 3) to ignore the presence of the other participants, but still change the playing strategy due to active exploration of the environment. And 4) there is also the possibility that behaviour was not altered consciously, but there were implicit biases at play. I will now explore each of these points in turn.

The first point on its own can be partially refuted because the behaviour during the social conditions was very different from the baseline, so we can conclude that the social conditions were not ignored (by conscious or implicit mechanisms). If not due to social factors differences from baseline can be a consequence of exploration. The baseline block was always first, hence, exploration may have started during the social blocks. Reduced risk-taking during the both-high condition could be an effect of exploration independent of social condition, which would be a confounding factor for our theory-driven conclusions. So, is it exploration or a specific strategy during the social conditions that drive different patterns from those seen during baseline? The answer to this question is partially revealed by the performance change between baseline and the social conditions, and from our additional exploratory analysis of individual differences in anxiety levels.

By dividing the participants into subsets of high and low anxiety levels we found that exploration alone may not be able to explain the results, at least not for participants with high anxiety levels who had similar behaviour to baseline, and more in line with the theory. Moreover, if exploration was higher during the social conditions we would expect fewer ‘correct’ answers (Sims et al., 2013). Choosing a stimulus with a higher EV from two stimuli with different EVs is considered ‘rational behaviour’ as it maximises the outcomes fully. Contrary, if suboptimal stimuli were chosen, this would mean exploration. We did not find indications of exploratory behaviour during the beginning of the social blocks compared to the final stages (measured by performance), and no difference between performance during baseline and during the social blocks. Meaning that the risk-taking patterns were not the result of exploration, but a response specific to the uncertainty of the different stimuli. This would not hold true if there was exploration employed during trials of the same means compared to trials of different means. For example, the participants may have been interested in exploring the high-variance stimulus only when it was paired with a stimulus of the same mean. Given the high entropy of our stimuli, we can assume that this highly demanding cognitive task that involves tracking pairs of stimuli and adjusting behavioural strategies according to the

different pairs is unlikely. Hence, we can conclude that participants were explicitly or implicitly more cautious at the beginning of the blocks during the presence of others. This explanation fits with the results we found for the participants with high anxiety, whose responses were more consistent with the baseline condition. This may have been due to conformity or resolving to model free behaviour which I will elaborate on later.

We cannot know if the participants were explicitly trying to benefit or hinder the in-group or out-group members, however, either the participants had a strategy to show more cautious behaviour during both social conditions (resulting in an initial drop in predicted risk-taking), or they were influenced by social factors such as increased stress or arousal while being observed. If we assume that the mere presence of in-group and out-group members affected the participants differentially and without their realisation (implicitly), we can interpret the results as a consequence of a reduced uncertainty in the in-group condition which led to the initial risk-taking or due to increased arousal and stress during social observation. A few theories can explain the corrected behaviour and the observed reduction in overall risk-taking during social conditions.

The initially increased risk-taking in the in-group condition, and general risk-averse behaviour in the initial stage of the out-group condition, are in line with the social identity theory (Cruwys et al., 2021). However, our prediction that risk-taking would increase overall in the in-group compared to the out-group condition was not supported. Although some studies have found that the mere presence of peers is enough to increase risk-taking (Chou and Nordgren, 2017) and neural responses to feedback (Hobson and Inzlicht, 2016), other studies have shown that the mere presence without input from peers does not lead to increased risk-taking, and even oppositely in young adults, can lead to risk-averse behaviours (Somerville et al., 2019). A recent meta-analysis suggested that mere observation by peers in adolescence (10 – 19 years of age) can drive risk-taking, however, this effect was very small unless there was an expression of risk-seeking preferences by peers (Powers et al., 2022). Consistently it has been suggested that risky behaviour is not influenced by being observed unless there are

certain beliefs about the observer's norms and values (Stone and Allgaier, 2008). As such, the behaviour of our participants may be a consequence of what they believed was best 'to show' to the observers. Although we did not find increased risk-taking under in-group observation, future studies should not only test different levels of peer engagement with young adults but also the influence of beliefs about the observer's values.

Additionally, our results did not support the hypothesis that the in-group and out-group conditions would lead to a different neural response at the time of outcome for wins and losses (Hobson and Inzlicht, 2016). These results may be due to wins and losses being defined as rewards higher (wins) or lower (losses) than a reference point, rather than rewarding and aversive events. Our findings can also suggest that the online presence of the observers was not enough to induce different neural responses between the conditions. Another possibility is that the neural response coming from multiple sources within the midbrain may become more distributed and more difficult to detect with EEG. Future studies should explore these possibilities and determine the influence of in- and out-group members on feedback evaluation.

4.4.2 Risk-taking under stress, arousal, and anxiety

Another explanation of our results can be general arousal due to the presence of others. Previously it has been shown that the mere presence of peers in adolescence is enough to induce a self-conscious emotional response, measured by skin conductance and paralleled by mPFC activation (Somerville et al., 2013). A similar arousal or stress response may have led to our results via increased feelings of uncertainty and decreased risk-taking (Mullen et al., 1997). Although the participants showed an initial risk-averse behaviour across the social blocks, separating different anxiety levels revealed an unexpected pattern. In the current work, we included trait and state anxiety measures with mostly exploratory aims. We aimed to test whether and how decision-making would be influenced by anxiety. Our results showed that participants with higher anxiety had

a behaviour that was more consistent with the theory we tested compared to those with low anxiety levels. High anxiety (state and trait) was positively correlated with risk-taking in the both-high condition during the in-group context. For the out-group condition, higher anxiety was correlated with more risk-taking in the both-high trials and less risk-taking in the both-low trials. This is contrary to the expectation that higher stress or anxiety would induce risk-averse behaviours across the task. Such a pattern may be explained by both the theory of risk-taking in the basal ganglia and the influence of stress and anxiety on decision-making.

Previously it has been suggested that anxiety and stress can lead to habit-based (model-free) decision-making (Porcelli and Delgado, 2017), which relies more on well-learned associations as opposed to making goal-based decisions (model-based, Dolan and Dayan, 2013). The observation that the participants with high anxiety were slightly more consistent with baseline in terms of behaviour may be due to goal-based exploration by low-anxiety participants and more reliance on model-free choices in high-anxiety ones. In terms of brain response, as model-based decisions under the theory we tested would mean a relationship between the neural response to stimulus-PEs and subsequent risk-taking, the lack of this relationship may suggest resolving to model-free decisions. The neural response we found for the stimulus-PE was similar across the baseline and the two social blocks. However, although Y amplitudes were predictive of risk-taking during baseline, those obtained from the two social conditions could not predict risk-taking behaviour. This suggests a weaker link between the EEG components for the stimulus-PE and behaviour during the social blocks. This could be a consequence of reduced prefrontal activity (due to stress) and relying more on basal ganglia ‘learned’ responses. As the signal from midbrain areas that reaches the scalp is lower in intensity and is underweighted in EEG, if the activity is mainly coming from the midbrain, we can expect to see lower activity in the EEG response. However, a lower response may also be the consequence of medial prefrontal areas interfering with each other, which has been observed before in EEG (Hauser et al., 2015). Previous studies have also shown that anxiety may influence learning from extreme PEs and may result in

improved learning rates in some experiments (Piray and Daw, 2021). Hence, better learning in the basal ganglia and less influence from prefrontal biases may have led to slightly improved performance in anxious individuals during social conditions. Future studies with fMRI and skin conductance should explore whether social stress from observation is indeed the factor influencing model-free strategies, mediated by anxiety and whether model-free decisions show reduced activation in prefrontal areas or employ a different set of areas within the PFC that may interfere with each other as measured by EEG.

Another interpretation is that participants with high anxiety did not want to explore possible options because they wanted to ‘show’ the optimal behaviour to the observers, independent if they were an in-group or an out-group member. Conformity is a key coping mechanism of individuals with general anxiety (Lewis et al., 20088), which may lead to acting according to ‘what is right’ rather than exploring potentially suboptimal options. Future studies should test this possibility by qualitative exploration of the consciously employed strategies by the participants, or by manipulating the level of relationship and involvement of the confederates in the experiments. It is possible that conformity will differ depending on the strength of the group membership (for example, friends compared to family or strangers).

Overall, our findings highlight the importance of testing individual differences and neurodiversity in value-based decision-making within a social context. We suggest that future studies can benefit from incorporating anxiety and arousal measures into learning and decision-making experiments.

4.4.3 Limitations and future directions

One limitation of the current design was that the observation of in-group and out-group participants was not in person but done online. This was a constraint we could not resolve because the size of our experimental EEG booth would not allow for another participant to be inside, and additionally because of constraints due to social distancing (due to the COVID-19 pandemic). Although we did not want to distract the participants by reminding them that someone was watching, future research implementing ‘online’ observation should make sure there are cues and reminders that the participant is being watched by in-group or out-group observers, or there should be an addition to the design that involves others in more than mere observation.

Another limitation related to our models was that we did not test dynamic learning rates. Niv et al. (2012) suggest that dynamic learning rates may predict behaviour better over trials. We resolved to fit the models to the full experiment because neither of the models fit well the results from the individual social conditions. The lack of good model fits for the social conditions separately may have been due to the lack of dynamic learning rates, and no modelling of exploration. Additionally, the lack of difference found between the two social conditions in the EEG component response may be due to interference from more deep brain areas that potentially led to a low EEG signal, rather than differences in brain sources. Although we set 50% accuracy of the classifier to be the threshold for significance, most of our results are under 60% accuracy. This could be due to the low signal-to-noise ratio in the EEG, or due to the limitations of EEG to detect reliable activity if the source is deep within the brain (Burle et al., 2015). As the theory mainly focuses on striatal and cortico-striatal connections, fMRI and specifically EEG-informed fMRI may be better suited to study stimulus-PE-driven risk-taking.

In summary, this chapter identified the risk-taking behaviour of participants while being observed by an in-group and an out-group observer. Additionally, we presented EEG results showing that the temporal EEG responses are similar between the two social conditions and comparable to the baseline. We found that anxiety can be a modulator of altered risk-taking during social conditions. We suggested that future studies should focus on determining whether these effects were due to exploration, induced stress during the social conditions, conformity that tilted behaviour towards the baseline, or implicit biases driven by the presence of others. Given our behavioural and brain results, we suggest that our findings are likely due to stress due to observation combined with anxiety, which potentially led to model-free decision-making in some participants. Although this study failed to replicate previous findings, it informs the current literature on decision-making under uncertainty, answering important questions about the mere presence of others and the degree of involvement of social factors necessary to induce significant behavioural change.

General Discussion

5.1 Risk-processing and risk-seeking in social contexts

With the current work, we employed computational methods and EEG recordings to test state-of-the-art theories of risk processing and risk-taking in relation to social decision-making. We showed a discrimination between risk, signed and unsigned risk prediction errors in the EEG signal, and found that social stimuli result in the initiation of spatiotemporal EEG components that are comparable to previous non-social tasks. We also supported a novel theory suggesting that the function of dopamine as a response-invigorating signal may interfere with its risk-promoting function, leading to higher risk-taking following exceeded expectations. We found a centro-frontal EEG component that was linked to the prediction error occurring after cue onset. Additionally, under the observation of confederates, we observed intriguing results of partially reduced initial risk-seeking that converged to the expected behaviour only after the middle of each block. And we showed that the stimulus-PE EEG component was attenuated during the social conditions and the relationship between this EEG component and risk-seeking was no longer present during the social conditions compared

to the baseline. These results are highly significant in the context of recent studies focusing on individual differences, stress, and mental health conditions as contributors to a dynamic decision-making response switching between goal-oriented, model-based, or habitual, model-free behaviours.

5.2 Encoding predicted risk separately from risk prediction errors

In our first experiment, we used facial stimuli as cues for rewards and integrated them into a task based on a previous electrophysiological study with rhesus monkeys (O'Neill and Schultz, 2013). We developed the task further to allow the separation between risk and signed and unsigned risk prediction errors. The findings were consistent with previous fMRI and EEG studies showing the separation between the three signals (Lauffs et al., 2020; Preuschoff et al., 2008).

By keeping average rewards constant, we suggested that risk is coded separately from rewards. It is important to note that this does not exclude the possibility that there may be an overlap between the coding of risks and rewards in some areas, and even by the same neurones. The animal studies we based our task design on, for example, O'Neill and Schultz (2010a), found that there is a large proportion of neurones that code risk, but not rewards, although there was also an overlap in some neurones that coded both. Additionally, according to the mean-variance approach that we tested, the variance needs to be coded separately from the expected value. According to other theories, like utility theory and prospect theory, the risk does not need to be coded separately. It can be inferred from the non-linearity of the utility function (Liu and Polak, 2007) or from the nonlinearity and the sign (for example, risk-seeking in the loss domain), as in prospect theory (Kahneman and Tversky, 1979). The nonlinearity of reward encoding was also confirmed to be related to differential activation of the

striatum for wins compared to losses (Tom et al., 2007). So, indeed under some circumstances, it may not be necessary to encode risks separate from rewards. But under some circumstances, risk-coding and risk-prediction errors become necessary for learning (Bossaerts, 2010). In stochastic environments, a statistically efficient learner (that is proposed to be the Bayesian brain) must estimate the noise of the environment alongside the action values or expected values (Piray and Daw, 2021). Under uncertainty, prediction errors interact with risk to drive learning. For example, captured by the Kalman filter and its estimation of volatility, when prediction risk is lower, it is expected that errors are informative and the learning rate is higher, whereas if prediction risk is high, rewards are too noisy to be meaningful and the learning rate adapts to be low (Soltani and Izquierdo, 2019). Oppositely, stochasticity captures short-term changes in the environment and advances learning rates when the fluctuations are higher, similar to the notion of how prediction errors work under RW (Nassar et al., 2010). Hence, risk, as an additional statistical moment, provides additional information about the dynamic changes in distributions, relevant for future short- and long-term predictions and adaptive learning.

As a learning signal, the response to risk and risk prediction errors may require similar components to those found previously for rewards. Further on the temporal profile of the components we found, the literature on reward prediction errors has identified a two-component response to rewards, which consists of salience, or the magnitude of the error, and value, which conveys further meaning about the error (for a review see Fouragnan et al., 2017; Schultz, 2016b). We found evidence for unsigned RiPE, which is equivalent to the magnitude of the error in risk predictions, to peak after the current risk, and to continue decreasing until the end of the 800-ms period after stimulus onset. On the other hand, signed RiPE which indicates if the error has a positive or negative sign, consisted of multiple peaks that continued beyond the end of the tested period after stimulus onset. The results of both RiPEs components are consistent with studies on dopamine showing a more sustained signal for uncertainty coding compared to rewards (Fiorillo et al., 2003). Whereas the differential profile of the two is consistent

with recent EEG findings on RiPEs (Lauffs et al., 2020). Although the signed and unsigned RiPE may signal the valence of the errors, and their magnitude, respectively, it is still unanswered whether they occur in a sequence or simultaneously. This may be determined by examining the components we found in the context of a learning task. As suggested previously signed and unsigned prediction errors (of rewards or risks) may be relevant for different stages of learning (Rouhani and Niv, 2021). Future studies should explore how risk prediction errors contribute to learning, and if similar to rewards the two types of errors influence trial-wise versus wider-scale learning (Diederer and Fletcher, 2021).

Our findings are also novel in showing generalisability to the social domain. The faces we chose were mapped to numerical rewards in a way that was meaningful and subjective for each participant and depended on trustworthiness ratings. Currently, there is an increased interest in delineating, first, the generalisability of risk processes across stimuli of different domains, including financial and monetary rewards, as well as perceptual or more naturalistic uncertainty (Loued-Khenissi et al., 2020; Schonberg et al., 2011), and second, the amount of overlap of neural characteristics when it comes to non-social compared to social decision-making (Rilling and Sanfey, 2011; Ruff and Fehr, 2014). Although our results are insufficient to address naturalistic social decisions or risk that is specifically related to social decision-making, they point towards a common risk processing that is generalised across domains. Previous findings have pointed out the additional influence of facial trustworthiness on risk tolerance in the gain and loss domains (Qi et al., 2022). Therefore, it is essential to test whether faces would have an additional bias on the estimation of risk in tasks combining both social and non-social cues.

Our understanding of risk prediction errors is still very limited. However, risk processing is essential for learning in uncertain environments (d’Acromont et al., 2009; Piray and Daw, 2021; Preuschoff and Bossaerts, 2007), especially those that involve non-social rewards and social cues that can reduce uncertainty and enhance decisions. Our first

experiment confirmed separate risk components, showing distinct temporal and spatial profiles across the span of single trials. This finding suggests that risk prediction errors may be represented based on their valence and surprise, similar to findings from the literature on reward prediction errors.

5.3 Risk-seeking influenced by positive reward prediction errors

In our second experiment we harnessed computational modelling to test the relationship between two proposed functions of dopamine: its increase as a response to positive reward prediction errors and its influence on short-term risk-seeking. We tested multiple models based on reinforcement learning, prospect theory, and others. The model that best fit the results was the one based on risk-seeking as a consequence of a dopamine response to reward prediction errors and the addition of a free parameter. We supported these results with an EEG component that represented the discrimination between positive and negative stimulus prediction errors, and which predicted risk-taking on a trial-by-trial basis.

We discussed a discrepancy in the literature on risk-taking behaviour. Some studies showed increased risk-taking in gain contexts (Niv et al., 2012), while others proposed increased risk-taking for losses and risk-averse behaviour for wins (Kahneman and Tversky, 1979). Our results showed that risk-taking depended on the preceding prediction errors on each trial, which is consistent with the suggestion of a prefrontal bias on learned values within the Go and No-Go pathways of the basal ganglia. Differences between our and previous findings are likely due to the differences in the complexity of the stimuli in different experiments and the distributions of associated rewards. It is possible that there is a switch between utilising simpler models to solve an environment of a few rewards, compared to one with full distributions, which would require more

complex modelling (Moeller et al., 2021). Additionally, it is possible that prediction errors in a task with high entropy (full distributions) lead to different experiences across participants. Unexpected uncertainty can occur when the environment is subjectively interpreted as more volatile, which increases learning rates (Soltani and Izquierdo, 2019). And as mentioned earlier, learning rates can differ across participants who mistake stochasticity (small changes) for volatility by overestimating PEs, and those who have correct estimates (Piray and Daw, 2021). Subsequently, it is not unlikely that demands for learning about the environment change risk-seeking, that is the opportunity to reduce entropy and that this differs across participants. These demands may be the complexity of stimuli and time constraints requiring slow or fast decisions. It is possible that such switches between risk-probing to increase learning versus ‘playing safe’ to avoid threat may be the consequence of top-down cognitive control, resulting in fast switching between the two (Choi et al., 2022). We did not test whether there are differences between groups of participants in the models that best explained their behaviour, however, this is highly likely due to the evidence for individual differences that we observed and should be explored further.

Moreover, also related to the harnessing of different models as a response to task demands, it is essential to consider the neural correlates of risky decision-making within the context of tasks that require learning cue-reward contingencies on the go versus choosing between learned stimuli (Fitzgerald et al., 2010). Learning on the go requires working memory and potentially higher influence of frontal areas on the striatum (Ott and Nieder, 2019), which may be the reason previous studies have found different active areas of the prefrontal cortex and striatum during both types of tasks (Fitzgerald et al., 2010). These types of learning demands can lead to a separate contribution of prefrontal biases on habitual behaviour, which requires further tests.

In summary, our study contributes to the research on value-based decision-making by demonstrating that short-term changes in expected rewards can affect risky decision-making, as reflected by EEG patterns. While we cannot definitively conclude that increased phasic dopamine levels are the cause of the observed risk-taking, our findings suggest that it is the most likely explanation based on the models we tested. The findings of this study are relevant for pathological gambling and addiction-associated conditions, suggesting that different dopamine functions can interfere with each other to produce spontaneous and sometimes undesirable behaviour.

5.4 Presence of others and anxiety on risky decision-making

Within the same experiment that showed stimulus-PEs to be driving risk-seeking, we implemented two conditions aiming to study the effects of social identity theory on risky choices. The participants were observed by either an in-group or an out-group member assigned by a minimal grouping procedure. We found that being observed by others reduced risk-seeking for both social conditions. However, an additional analysis suggested that participants with higher anxiety did not change their risk-taking patterns compared to baseline as much as the participants with low anxiety.

Previously it has been suggested that uncertainty is reduced during the presence of in-group members, which subsequently increases risk-taking (Cruwys et al., 2021). The result of our study did not support these predictions from social identity theory. Rather, we suggested that the change in results from baseline may be due to increased exploration by participants with low levels of anxiety. Recently, it was found that trait anxiety can reduce exploration in different ways depending on relative versus total uncertainty (Fan et al., 2023). Fan et al. (ibid.) reported that anxious individuals were less likely to explore when perceiving high relative uncertainty. In our case, relative uncertainty would be deviating from what was learned during baseline, which is in line with the

consistency in the behaviour of participants with higher anxiety across the blocks. Future studies should test the relationship between risk-taking and anxiety in a causal manner by modulating the stochasticity and volatility of the rewards. Additionally, we suggested that arousal can be measured by pupillometry or skin conductance measures to determine if minimal stress can induce model-free decision-making as we suggested earlier (Leuchs et al., 2019).

Furthermore, unlike the baseline condition, our EEG results at stimulus-PE did not predict trial-by-trial risk-taking. We speculated that this finding may be due to the utilisation of more medial brain areas or a combination of prefrontal brain areas which would attenuate the EEG signal. We also did not show differential activity patterns at feedback-onset for wins and losses between the two social conditions as suggested previously (Hobson and Inzlicht, 2016). Potential explanations include failure of the membership bias due to the simple online observation which may not have been enough to induce ‘in-group’ and ‘out-group’ feelings, or the lack of actual wins and losses within our task. Future studies should test different strengths of the group membership bias by varying the involvement of others in the experimental tasks and exploring risk-taking under actual wins and losses.

Overall the social conditions did not result in most of the expected findings, however, this experiment adds important answers to the research of social neuroscience and proposes testable confounding factors in tasks with social observation. Notably, the results suggest that it is essential to measure exploration, and arousal, and to vary the level of social interaction in order to study risk-taking within social contexts. We further showed that it is important to measure individual differences or mental health and their interaction with social components to causally alter learning, risk-taking, decision-making, and the associated neural processes.

5.4.1 Implications for addiction and gambling

It would be unexpected to mention risk-seeking without mentioning the extremes which are addictions, substance abuse, gambling, and unwanted repetitive behaviours. The activation of the ventral medial prefrontal cortex, which is a crucial element in decision-making based on values, has been linked to addictive behaviours such as gambling, substance abuse, and OCD (Cavedini et al., 2002). A study also found that individuals with damage to the vmPFC performed worse on a task involving risky decision-making and showed more frequent and larger bets in a simulated gambling task, suggesting that the vmPFC plays a crucial role in decision-making related to risk and reward (Clark et al., 2008). The influence of the frontal cortex over the striatum has been recognised before in repetitive behaviours. In a study with mice, repetitive behaviour was found to be the result of impaired downregulation from the cortex to the ventral striatum, which was improved by optogenetic stimulation of the lateral OFC and its striatal terminals (Burguière et al., 2013). Although humans may be far from receiving such targeted therapy, the findings show that repetitive behaviours can be reversed by improving top-down regulation from the cortex. Additionally, reduced activation of the ventral striatum has been connected to the strong desire or urges experienced during gambling and substance addictions (Brewer and Potenza, 2008; Potenza, 2008). There is evidence on neurotransmitter levels that the overall increase in tonic dopamine in the striatum of medicated Parkinson's patients leads to increased gambling in about 13% of cases (Zhang et al., 2021). This tendency, however, is not seen in all patients and depends largely on personality characteristics and predispositions to gambling (Wolfschlag and Håkansson, 2023). As such, it is important to consider once again that the effects of dopamine are neuromodulatory and add a bias to a system, rather than definitively driving one behaviour or another. With the findings of our second experiment, we showed that fluctuations in prediction errors, confirmed by EEG components, are

sufficient to affect risk-seeking. As suggested by the theory we tested, this may be a consequence of transient dopamine change, affected by the influence of the prefrontal cortex on the striatum. These emerging findings present a captivating and promising frontier for scientific exploration.

5.5 Limitations and future directions

One of the limitations of the first study is the inability to directly compare social to non-social responses. Although we implemented a task and design that were directly comparable to animal studies, it is still unclear what would be the direct comparison between the risk and risk prediction errors coding between tasks with social and non-social stimuli. Moreover, one can argue that the social stimuli were ignored while the participants only encoded the numerical value associated with each face. The faces were specifically selected based on each participant's ratings of them according to a trust game. As such, the points mapped to each face were directly related to a social rating of trustworthiness. Although this suggests that there is a social component that takes place in the associations between the faces and numerical values, it is possible that this social component was completely ignored and the faces were simply used as reminders of the values of higher interest, namely, the numerical points. In empirical studies with human participants, there is a trade-off between the number of trials per condition that can be collected and the sustained attention of the participant throughout the task. Our experiment spanned over 40 minutes and involved no active choice, with about one-hour of preparation of the EEG device. Due to this, we were unable to implement a second condition with non-social stimuli and keep participants' attention sustained. Future studies could solve this issue by dividing the experiment into separate sessions or increasing the level of social involvement. For example, recent studies in reward PEs

have implemented a social paradigm in which PEs are tested as a response to observing another person's learning experience – termed simulated other's PEs (Suzuki et al., 2012, 2016). Future studies could implement a similar paradigm for risk processes and explore the simulated RiPEs while participants observe another person's behaviour.

For the second experiment, a similar limitation can be identified. Although this time we included a baseline condition prior to the two social conditions, as discussed earlier, this may have resulted in possible exploratory behaviour in some participants during the social blocks. Hence, some variation in risk-taking during the social conditions may be explained by the timeline of the task rather than the social biases that we predicted. Such exploration may especially be the result of the low stakes of the task. As the rewards to be won were points later converted into a potential small monetary reward, motivation to maximise the winnings may not be sufficient to reduce choosing suboptimal options. Future studies aiming to keep behaviour constant across multiple blocks should either implement larger rewards or include losses as well as rewards. Such manipulation may keep the focus on maximising the final pay-off and avoid unexpected exploratory behaviour.

5.6 Conclusion

With this work, we examined the neural and computational processes of social and non-social risky decision-making. We first measured the brain responses to risk and risk prediction errors during passive observation of the selection of risky stimuli. We found evidence for three components right after the presentation of the stimuli that were correlated with risk, signed and unsigned risk prediction errors. Next, we tested a recently proposed model that risk-seeking on a trial-by-trial basis is influenced by positive prediction errors, a bias that likely takes place within the basal ganglia and is related to the mechanisms of Go and No-Go pathways within cortico-striato-thalamo-cortical circuits. Our results supported the model and additionally showed the temporal

and spatial profile of EEG activity related to the stimulus prediction error. Additionally, we tested predictions from social identity theory and showed that the mere online presence of others led to different behavioural results from baseline, although more research is necessary to determine the exact factors contributing to this change.

This thesis presents a significant advancement in the fields of cognitive neuroscience and neuroeconomics by expanding upon the neural and computational foundations of decision-making in risky contexts, both social and non-social. The insights provided offer valuable perspectives for future research endeavours seeking to unravel the complex intricacies of behaviour and brain responses under uncertainty.

Bibliography

- Abbott, Larry F and Peter Dayan (1999). ‘The effect of correlated variability on the accuracy of a population code’. In: *Neural computation* 11.1, pp. 91–101.
- Abler, Birgit et al. (June 2006). ‘Prediction error as a linear function of reward probability is coded in human nucleus accumbens’. In: *NeuroImage* 31 (2), pp. 790–795. ISSN: 1053-8119. DOI: 10.1016/J.NEUROIMAGE.2006.01.001.
- Abrams, Dominic et al. (1990). ‘Knowing what to think by knowing who you are: Self-categorization and the nature of norm formation, conformity and group polarization’. In: *British Journal of Social Psychology* 29 (2). ISSN: 20448309. DOI: 10.1111/j.2044-8309.1990.tb00892.x.
- Abuz-Akel, Ahmad and Simone Shamay-Tsoory (2011). ‘Neuroanatomical and neurochemical bases of theory of mind’. In: *Neuropsychologia* 49 (11). ISSN: 00283932. DOI: 10.1016/j.neuropsychologia.2011.07.012.
- Akpan, Ben (2020). ‘Classical and Operant Conditioning—Ivan Pavlov; Burrhus Skinner BT - Science Education in Theory and Practice: An Introductory Guide to Learning Theory’. In: ed. by Ben Akpan and Teresa J Kennedy, pp. 71–84. DOI: 10.1007/978-3-030-43620-9_6.
- Albert, Dustin, Jason Chein and Laurence Steinberg (2013). ‘The Teenage Brain: Peer Influences on Adolescent Decision Making’. In: *Current Directions in Psychological Science* 22 (2). ISSN: 14678721. DOI: 10.1177/0963721412471347.
- Andersen, Steffen et al. (2014). ‘Estimating subjective probabilities’. In: *Journal of Risk and Uncertainty* 48, pp. 207–229.
- Arabadzhiyska, Desislava H. et al. (2021). ‘A common neural currency account for social and non-social decisions’. In: *bioRxiv*.

- Averbeck, Bruno B. and Brad Duchaine (2009). 'Integration of Social and Utilitarian Factors in Decision Making'. In: *Emotion* 9 (5). ISSN: 15283542. DOI: 10.1037/a0016509.
- Balasubramani, Pragathi P et al. (2014). 'An extended reinforcement learning model of basal ganglia to understand the contributions of serotonin and dopamine in risk-based decision making, reward prediction, and punishment learning'. In: *Frontiers in computational neuroscience* 8, p. 47. ISSN: 1662-5188.
- Ballew, Charles C. and Alexander Todorov (2007). 'Predicting political elections from rapid and unreflective face judgments'. In: *Proceedings of the National Academy of Sciences of the United States of America* 104 (46). ISSN: 00278424. DOI: 10.1073/pnas.0705435104.
- Balter, Lawrence and Catherine S. Tamis-LeMonda (2016). *Child psychology: A handbook of contemporary issues: Third edition*. DOI: 10.4324/9781315764931.
- Bargh, John A. et al. (2012). 'Automaticity in social-cognitive processes'. In: *Trends in Cognitive Sciences* 16 (12). ISSN: 13646613. DOI: 10.1016/j.tics.2012.10.002.
- Bateson, Melissa (2002). 'Recent advances in our understanding of risk-sensitive foraging preferences'. In: *Proceedings of the Nutrition Society* 61 (4), pp. 509–516. ISSN: 1475-2719.
- Bavel, Jay J. Van, Dominic J. Packer and William A. Cunningham (2008). 'The Neural Substrates of In-Group Bias'. In: *Psychological Science* 19 (11). ISSN: 0956-7976. DOI: 10.1111/j.1467-9280.2008.02214.x.
- Bechara, Antoine and Antonio R. Damasio (2005). 'The somatic marker hypothesis: A neural theory of economic decision'. In: *Games and Economic Behavior* 52 (2). ISSN: 10902473. DOI: 10.1016/j.geb.2004.06.010.
- Behrens, Timothy E.J. et al. (2007). 'Learning the value of information in an uncertain world'. In: *Nature Neuroscience* 10 (9). ISSN: 10976256. DOI: 10.1038/nn1954.
- Bellemare, Marc G., Will Dabney and Rémi Munos (2017). 'A distributional perspective on reinforcement learning'. In: vol. 1.

- Ben-Ner, Avner et al. (2009a). ‘Identity and in-group/out-group differentiation in work and giving behaviors: Experimental evidence’. In: *Journal of Economic Behavior & Organization* 72.1, pp. 153–170.
- Ben-Ner, Avner et al. (2009b). ‘Identity and in-group/out-group differentiation in work and giving behaviors: Experimental evidence’. In: *Journal of Economic Behavior and Organization* 72 (1). ISSN: 01672681. DOI: 10.1016/j.jebo.2009.05.007.
- Bernhardt, Boris C. and Tania Singer (2012). ‘The neural basis of empathy’. In: *Annual Review of Neuroscience* 35. ISSN: 0147006X. DOI: 10.1146/annurev-neuro-062111-150536.
- Bissonette, Gregory B. and Matthew R. Roesch (2016). ‘Neurophysiology of reward-guided behavior: Correlates related to predictions, value, motivation, errors, attention, and action’. In: *Current Topics in Behavioral Neurosciences* 27. ISSN: 18663389. DOI: 10.1007/7854_2015_382.
- Bogacz, Rafal (2017). ‘Theory of reinforcement learning and motivation in the basal ganglia’. In: *BioRxiv*, p. 174524.
- Bogacz, Rafal and Tobias Larsen (2011). ‘Integration of reinforcement learning and optimal decision-making theories of the basal ganglia’. In: *Neural computation* 23.4, pp. 817–851.
- Bortoletto, Marta, Marianna J. Lemonis and Ross Cunnington (2011). ‘The role of arousal in the preparation for voluntary movement’. In: *Biological Psychology* 87 (3). ISSN: 03010511. DOI: 10.1016/j.biopsycho.2011.04.008.
- Bossaerts, Peter (2010). ‘Risk and risk prediction error signals in anterior insula.’ In: *Brain structure function* 214 (5-6). ISSN: 18632661. DOI: 10.1007/s00429-010-0253-1.
- Brand, Matthias et al. (2007). ‘Role of the amygdala in decisions under ambiguity and decisions under risk: Evidence from patients with Urbach-Wiethe disease’. In: *Neuropsychologia* 45 (6). ISSN: 00283932. DOI: 10.1016/j.neuropsychologia.2006.09.021.

- Brewer, Judson A. and Marc N. Potenza (2008). 'The neurobiology and genetics of impulse control disorders: Relationships to drug addictions'. In: *Biochemical Pharmacology* 75 (1). ISSN: 00062952. DOI: 10.1016/j.bcp.2007.06.043.
- Brown, Darin R., Sarah Pirio Richardson and James F. Cavanagh (2020). 'An EEG marker of reward processing is diminished in Parkinson's disease'. In: *Brain Research* 1727. ISSN: 18726240. DOI: 10.1016/j.brainres.2019.146541.
- Brown, Joshua W. and Todd S. Braver (2007). 'Risk prediction and aversion by anterior cingulate cortex'. In: *Cognitive, Affective and Behavioral Neuroscience* 7 (4). ISSN: 15307026. DOI: 10.3758/CABN.7.4.266.
- Broyd, Samantha J. et al. (2012). 'An electrophysiological monetary incentive delay (e-MID) task: A way to decompose the different components of neural response to positive and negative monetary reinforcement'. In: *Journal of Neuroscience Methods* 209 (1). ISSN: 01650270. DOI: 10.1016/j.jneumeth.2012.05.015.
- Burguière, Eric et al. (2013). 'Optogenetic stimulation of lateral orbitofronto-striatal pathway suppresses compulsive behaviors'. In: *Science* 340.6137, pp. 1243–1246.
- Burke, Christopher J. and Philippe N. Tobler (2011). 'Reward skewness coding in the insula independent of probability and loss'. In: *Journal of Neurophysiology* 106 (5). ISSN: 00223077. DOI: 10.1152/jn.00471.2011.
- Burle, Borís et al. (2015). 'Spatial and temporal resolutions of EEG: Is it really black and white? A scalp current density view'. In: *International Journal of Psychophysiology* 97 (3). ISSN: 18727697. DOI: 10.1016/j.ijpsycho.2015.05.004.
- Bush, Robert R. and Frederick Mosteller (1951). 'A mathematical model for simple learning'. In: *Psychological Review* 58 (5). ISSN: 0033295X. DOI: 10.1037/h0054388.
- Caplin, Andrew and Paul W. Glimcher (2013). 'Basic Methods from Neoclassical Economics'. In: *Neuroeconomics: Decision Making and the Brain: Second Edition*. DOI: 10.1016/B978-0-12-416008-8.00001-2.
- Cartar, Ralph V (1991). 'A test of risk-sensitive foraging in wild bumble bees'. In: *Ecology* 72 (3), pp. 888–895. ISSN: 1939-9170.

- Cavedini, Paolo et al. (2002). 'Frontal lobe dysfunction in pathological gambling patients'. In: *Biological Psychiatry* 51 (4). ISSN: 00063223. DOI: 10.1016/S0006-3223(01)01227-6.
- Charles, D. Spielberger et al. (1972). 'State-Trait Anxiety Inventory for Adults'. In.
- Chaua, Bolton K.H. et al. (2018). 'Dopamine and reward: A view from the prefrontal cortex'. In: *Behavioural Pharmacology* 29 (7). ISSN: 14735849. DOI: 10.1097/FBP.000000000000424.
- Choi, Eun A et al. (2022). 'A corticothalamic circuit trades off speed for safety during decision-making under motivational conflict'. In: *Journal of Neuroscience* 42.16, pp. 3473–3483.
- Chou, Eileen Y. and Loran F. Nordgren (2017). 'Safety in Numbers: Why the Mere Physical Presence of Others Affects Risk-taking Behaviors'. In: *Journal of Behavioral Decision Making* 30 (3). ISSN: 10990771. DOI: 10.1002/bdm.1959.
- Cikara, Mina, Matthew M. Botvinick and Susan T. Fiske (2011). 'Us versus them: Social identity shapes neural responses to intergroup competition and harm'. In: *Psychological Science* 22 (3). ISSN: 09567976. DOI: 10.1177/0956797610397667.
- Ciranka, Simon and Wouter van den Bos (2021). 'Adolescent risk-taking in the context of exploration and social influence'. In: *Developmental Review* 61. ISSN: 02732297. DOI: 10.1016/j.dr.2021.100979.
- Clark, L et al. (2008). 'Differential effects of insular and ventromedial prefrontal cortex lesions on risky decision-making'. In: *Brain* 131 (5), pp. 1311–1322. ISSN: 1460-2156.
- Cobb, Roger W and David M Primo (2004). *The plane truth: Airline crashes, the media, and transportation policy*. Rowman & Littlefield.
- Corlett, Philip R., Jessica A. Mollick and Hedy Kober (2022). 'Meta-analysis of human prediction error for incentives, perception, cognition, and action'. In: *Neuropsychopharmacology* 47 (7). ISSN: 1740634X. DOI: 10.1038/s41386-021-01264-3.
- Corradi-Dell'Acqua, Corrado et al. (2016). 'Cross-modal representations of first-hand and vicarious pain, disgust and fairness in insular and cingulate cortex'. In: *Nature Communications* 7. ISSN: 20411723. DOI: 10.1038/ncomms10904.

- Craver, Carl F. (2005). 'Beyond reduction: Mechanisms, multifield integration and the unity of neuroscience'. In: *Studies in History and Philosophy of Science Part C :Studies in History and Philosophy of Biological and Biomedical Sciences* 36 (2 SPEC. ISS.). ISSN: 13698486. DOI: 10.1016/j.shpsc.2005.03.008.
- Critchley, Hugo D., Christopher J. Mathias and Raymond J. Dolan (2001). 'Neural activity in the human brain relating to uncertainty and arousal during anticipation'. In: *Neuron* 29 (2). ISSN: 08966273. DOI: 10.1016/S0896-6273(01)00225-2.
- Cruwys, Tegan, Mark Stevens and Katharine H. Greenaway (2020a). 'A social identity perspective on COVID-19: Health risk is affected by shared group membership'. In: *British Journal of Social Psychology* 59 (3). ISSN: 20448309. DOI: 10.1111/bjso.12391.
- Cruwys, Tegan et al. (2020b). 'Risk-Taking That Signals Trust Increases Social Identification'. In: *Social Psychology* 51 (5). ISSN: 21512590. DOI: 10.1027/1864-9335/a000417.
- Cruwys, Tegan et al. (2021). 'When trust goes wrong: A social identity model of risk taking'. In: *Journal of Personality and Social Psychology* 120 (1). ISSN: 00223514. DOI: 10.1037/pspi0000243.
- Dabney, Will et al. (2020). 'A distributional code for value in dopamine-based reinforcement learning'. In: *Nature* 577.7792, pp. 671–675.
- d'Acromont, Mathieu et al. (2009). 'Neural correlates of risk prediction error during reinforcement learning in humans'. In: *Neuroimage* 47 (4), pp. 1929–1939. ISSN: 1053-8119.
- Dahl, Ronald E. (2004). 'Adolescent brain development: A period of vulnerabilities and opportunities - Keynote Address'. In: vol. 1021. DOI: 10.1196/annals.1308.001.
- Daunizeau, Jean (n.d.). 'MBB-team/VBA-toolbox'. In: (). Accessed: 2022-10-03.
- Decety, Jean (2010). 'The neurodevelopment of empathy in humans'. In: *Developmental Neuroscience* 32 (4). ISSN: 03785866. DOI: 10.1159/000317771.
- Decker, Johannes H. et al. (2016). 'From Creatures of Habit to Goal-Directed Learners: Tracking the Developmental Emergence of Model-Based Reinforcement Learning'. In: *Psychological Science* 27 (6). ISSN: 14679280. DOI: 10.1177/0956797616639301.

- Defoe, Ivy N. et al. (2015). 'A meta-analysis on age differences in risky decision making: Adolescents versus children and adults'. In: *Psychological Bulletin* 141 (1). ISSN: 00332909. DOI: 10.1037/a0038088.
- Delorme, Arnaud and Scott Makeig (2004). 'EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis'. In: *Journal of Neuroscience Methods* 134 (1). ISSN: 01650270. DOI: 10.1016/j.jneumeth.2003.10.009.
- Deng, Leyou et al. (2023). 'Distinct neural dynamics underlying risk and ambiguity during valued-based decision making'. In: *Psychophysiology* 60 (3). ISSN: 14698986. DOI: 10.1111/psyp.14201.
- Diederer, Kelly MJ and Paul C Fletcher (2021). 'Dopamine, prediction error and beyond'. In: *The Neuroscientist* 27.1, pp. 30–46.
- Dodd, M. Leann et al. (2005). 'Pathological gambling caused by drugs used to treat Parkinson disease'. In: *Archives of Neurology* 62 (9). ISSN: 00039942. DOI: 10.1001/archneur.62.9.noc50009.
- Dolan, Ray J. and Peter Dayan (2013). 'Goals and habits in the brain'. In: *Neuron* 80 (2). ISSN: 08966273. DOI: 10.1016/j.neuron.2013.09.007.
- Dreu, Carsten K.W. De et al. (2010). 'The neuropeptide oxytocin regulates parochial altruism in intergroup conflict among humans'. In: *Science* 328 (5984). ISSN: 00368075. DOI: 10.1126/science.1189047.
- Driver-Dunckley, E., J. Samanta and Mark Stacy (2003). 'Pathological gambling associated with dopamine agonist therapy in Parkinson's disease'. In: *Neurology* 61 (3). ISSN: 00283878. DOI: 10.1212/01.WNL.0000076478.45005.EC.
- Dunbar, Robin I.M. (1998). 'The social brain hypothesis'. In: *Evolutionary Anthropology* 6 (5). ISSN: 10601538. DOI: 10.1002/(SICI)1520-6505(1998)6:5<178::AID-EVAN5>3.0.CO;2-8.
- d'Acremont, Mathieu and Peter Bossaerts (2008). 'Neurobiological studies of risk assessment: a comparison of expected utility and mean-variance approaches'. In: *Cognitive, Affective, Behavioral Neuroscience* 8 (4), pp. 363–374. ISSN: 1531-135X.

- Eckel, Catherine C. and Rick K. Wilson (2003). ‘The Human face of game theory: Trust and reciprocity in sequential games’. In: *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research* 9781610444347.
- Elliott, Rebecca and Bill Deakin (2005). ‘Role of the Orbitofrontal Cortex in Reinforcement Processing and Inhibitory Control: Evidence from functional magnetic resonance imaging Studies in Healthy Human Subjects’. In: *International Review of Neurobiology* 65. ISSN: 00747742. DOI: 10.1016/S0074-7742(04)65004-5.
- Fan, Haoxue, Samuel J Gershman and Elizabeth A Phelps (2023). ‘Trait somatic anxiety is associated with reduced directed exploration and underestimation of uncertainty’. In: *Nature Human Behaviour* 7.1, pp. 102–113.
- Fazio, Lisa K. and Elizabeth J. Marsh (2009). ‘Surprising feedback improves later memory’. In: *Psychonomic Bulletin and Review* 16 (1). ISSN: 10699384. DOI: 10.3758/PBR.16.1.88.
- Fazl, Arash and Jori Fleisher (2018). ‘Anatomy, physiology, and clinical syndromes of the basal ganglia: a brief review’. In: *Seminars in pediatric neurology*. Vol. 25. Elsevier, pp. 2–9.
- Ferrari-Toniolo, Simone and Wolfram Schultz (2021). ‘Reliable population utility signal from diverse economic value coding in orbitofrontal neurons’. In: *bioRxiv*.
- Ferry, Amon T. et al. (2000). ‘Prefrontal cortical projections to the striatum in macaque monkeys: Evidence for an organization related to prefrontal networks’. In: *Journal of Comparative Neurology* 425 (3). ISSN: 00219967. DOI: 10.1002/1096-9861(20000925)425:3<447::AID-CNE9>3.0.CO;2-V.
- Finucane, Melissa L. et al. (2000). ‘The affect heuristic in judgments of risks and benefits’. In: *Journal of Behavioral Decision Making* 13 (1). ISSN: 08943257. DOI: 10.1002/(SICI)1099-0771(200001/03)13:1<1::AID-BDM333>3.0.CO;2-S.
- Fiorillo, Christopher D (2011). ‘Transient activation of midbrain dopamine neurons by reward risk’. In: *Neuroscience* 197, pp. 162–171.
- Fiorillo, Christopher D, Philippe N Tobler and Wolfram Schultz (2003). ‘Discrete coding of reward probability and uncertainty by dopamine neurons’. In: *Science* 299 (5614), pp. 1898–1902. ISSN: 0036-8075.

- Fiorillo, Christopher D, Philippe N Tobler and Wolfram Schultz (2005). ‘Evidence that the delay-period activity of dopamine neurons corresponds to reward uncertainty rather than backpropagating TD errors’. In: *Behavioral and brain Functions* 1 (1), pp. 1–5. ISSN: 1744-9081.
- Fitzgerald, Thomas H B et al. (2010). ‘Differentiable Neural Substrates for Learned and Described Value and Risk’. In: *Current Biology* 20.
- Fouragnan, Elsa, Chris Retzler and Marios G. Philiastides (July 2018). ‘Separate neural representations of prediction error valence and surprise: Evidence from an fMRI meta-analysis’. In: *Human Brain Mapping* 39 (7), pp. 2887–2906. ISSN: 1097-0193. DOI: 10.1002/HBM.24047.
- Fouragnan, Elsa et al. (2015). ‘Two spatiotemporally distinct value systems shape reward-based learning in the human brain’. In: *Nature Communications* 6. ISSN: 20411723. DOI: 10.1038/ncomms9107.
- Fouragnan, Elsa et al. (2017). ‘Spatiotemporal neural characterization of prediction error valence and surprise during reward learning in humans’. In: *Scientific reports* 7 (1), pp. 1–18. ISSN: 2045-2322.
- Frank, Michael J. et al. (2007). ‘Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning’. In: *Proceedings of the National Academy of Sciences of the United States of America* 104 (41). ISSN: 00278424. DOI: 10.1073/pnas.0706111104.
- Frank, Michael J. et al. (2009). ‘Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation’. In: *Nature Neuroscience* 12 (8). ISSN: 10976256. DOI: 10.1038/nn.2342.
- Franzen, Léon et al. (2020). ‘Auditory information enhances post-sensory visual evidence during rapid multisensory decision-making’. In: *Nature Communications* 11 (1). ISSN: 20411723. DOI: 10.1038/s41467-020-19306-7.
- Fu, Hong et al. (2022). ‘Subjective and objective risk perceptions and the willingness to pay for agricultural insurance: evidence from an in-the-field choice experiment in rural China’. In: *The Geneva Risk and Insurance Review*, pp. 1–24.

- Fujino, Junya et al. (2020). 'Role of the right temporoparietal junction in intergroup bias in trust decisions'. In: *Human Brain Mapping* 41 (6). ISSN: 10970193. DOI: 10.1002/hbm.24903.
- Fukunaga, Rena, John R. Purcell and Joshua W. Brown (2018). 'Discriminating formal representations of risk in anterior cingulate cortex and inferior frontal gyrus'. In: *Frontiers in Neuroscience* 12 (AUG). ISSN: 1662453X. DOI: 10.3389/fnins.2018.00553.
- Gardner, Margo and Laurence Steinberg (2005). 'Peer influence on risk taking, risk preference, and risky decision making in adolescence and adulthood: An experimental study'. In: *Developmental Psychology* 41 (4). ISSN: 00121649. DOI: 10.1037/0012-1649.41.4.625.
- Gerfen, Charles R. and W. Scott Young (1988). 'Distribution of striatonigral and striatopallidal peptidergic neurons in both patch and matrix compartments: an in situ hybridization histochemistry and fluorescent retrograde tracing study'. In: *Brain Research* 460 (1). ISSN: 00068993. DOI: 10.1016/0006-8993(88)91217-6.
- Gherman, Sabina and Marios G Philiastides (2014). 'Neural representations of confidence emerge from the process of decision formation during perceptual choices'. In: DOI: 10.1016/j.neuroimage.2014.11.036.
- Gilbertson, Tom, David Arkadir and J. Douglas Steele (2020). 'Opposing patterns of abnormal D1 and D2 receptor dependent cortico-striatal plasticity explain increased risk taking in patients with DYT1 dystonia'. In: *PLoS ONE* 15 (5). ISSN: 19326203. DOI: 10.1371/journal.pone.0226790.
- Gilbertson, Tom and Douglas Steele (2021). 'Tonic dopamine, uncertainty and basal ganglia action selection'. In: *Neuroscience* 466. ISSN: 18737544. DOI: 10.1016/j.neuroscience.2021.05.010.
- Giorgetta, Cinzia et al. (2012). 'Reduced risk-taking behavior as a trait feature of anxiety'. In: *Emotion* 12 (6). ISSN: 15283542. DOI: 10.1037/a0029119.

- Glimcher, Paul W. (2011). 'Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis (Proceedings of the National Academy of Sciences of the United States of America (2011) 108, S3, (15647-15654) DOI: 10.1073/pnas.1014269108)'. In: *Proceedings of the National Academy of Sciences of the United States of America* 108 (42). ISSN: 00278424. DOI: 10.1073/pnas.1114363108.
- Gold, Joshua I and Michael N Shadlen (2007). 'The neural basis of decision making'. In: *Annual review of neuroscience* 30.
- Goldstein, Rita Z. et al. (2006). 'The effect of graded monetary reward on cognitive event-related potentials and behavior in young healthy adults'. In: *International Journal of Psychophysiology* 62 (2). ISSN: 01678760. DOI: 10.1016/j.ijpsycho.2006.05.006.
- Grabenhorst, Fabian et al. (2019). 'Primate prefrontal neurons signal economic risk derived from the statistics of recent reward experience'. In: *eLife* 8. ISSN: 2050084X. DOI: 10.7554/eLife.44838.
- Greening, Steven et al. (2014). 'Individual differences in the anterior insula are associated with the likelihood of financially helping versus harming others'. In: *Cognitive, Affective and Behavioral Neuroscience* 14 (1). ISSN: 15307026. DOI: 10.3758/s13415-013-0213-3.
- Grima, Laura L. et al. (2022). 'Nucleus accumbens D1-receptors regulate and focus transitions to reward-seeking action'. In: *Neuropsychopharmacology* 47 (9). ISSN: 1740634X. DOI: 10.1038/s41386-022-01312-6.
- Grossman, Cooper D., Bilal A. Bari and Jeremiah Y. Cohen (2022). 'Serotonin neurons modulate learning rate through uncertainty'. In: *Current Biology* 32 (3). ISSN: 18790445. DOI: 10.1016/j.cub.2021.12.006.
- Guan, Maime et al. (2020). 'A cognitive modeling analysis of risk in sequential choice tasks'. In: *Judgment and Decision Making* 15 (5). ISSN: 19302975. DOI: 10.1017/s1930297500007956.

- Guitart-Masip, Marc et al. (2011). 'Vigor in the face of fluctuating rates of reward: An experimental examination'. In: *Journal of Cognitive Neuroscience* 23 (12). ISSN: 0898929X. DOI: 10.1162/jocn_a_00090.
- Hare, Todd A. et al. (May 2008). 'Dissociating the Role of the Orbitofrontal Cortex and the Striatum in the Computation of Goal Values and Prediction Errors'. In: *Journal of Neuroscience* 28 (22), pp. 5623–5630. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.1309-08.2008. URL: <https://www.jneurosci.org/content/28/22/5623>.
- Harris, Lasana T. and Susan T. Fiske (2010). 'Neural regions that underlie reinforcement learning are also active for social expectancy violations'. In: *Social Neuroscience* 5 (1). ISSN: 17470919. DOI: 10.1080/17470910903135825.
- Hauser, Tobias U. et al. (2015). 'Temporally dissociable contributions of human medial prefrontal subregions to reward-guided learning'. In: *Journal of Neuroscience* 35 (32). ISSN: 15292401. DOI: 10.1523/JNEUROSCI.0560-15.2015.
- Hecht, Stephen S (2006). 'Cigarette smoking: cancer risks, carcinogens, and mechanisms'. In: *Langenbeck's archives of surgery* 391, pp. 603–613.
- Heekeren, Hauke R., Sean Marrett and Leslie G. Ungerleider (2008). 'The neural systems that mediate human perceptual decision making'. In: *Nature Reviews Neuroscience* 9 (6). ISSN: 1471003X. DOI: 10.1038/nrn2374.
- Helversen, Bettina von et al. (2018). 'Foraging, exploration, or search? On the (lack of) convergent validity between three behavioral paradigms'. In: *Evolutionary Behavioral Sciences* 12 (3). ISSN: 23302933. DOI: 10.1037/ebs0000121.
- Hertel, Guido, Norbert L. Kerr and Lawrence A. Messé (2000). 'Motivation gains in performance groups: Paradigmatic and theoretical developments on the Köhler effect'. In: *Journal of Personality and Social Psychology* 79 (4). ISSN: 00223514. DOI: 10.1037/0022-3514.79.4.580.
- Hirschbichler, Stephanie T, John C Rothwell and Sanjay G Manohar (2022). 'Dopamine increases risky choice while D2 blockade shortens decision time'. In: *Experimental Brain Research* 240, pp. 3351–3360. DOI: 10.1007/s00221-022-06501-9. URL: <https://doi.org/10.1007/s00221-022-06501-9>.

- Hobson, Nicholas M. and Michael Inzlicht (2016). 'The mere presence of an outgroup member disrupts the brain's feedback-monitoring system'. In: *Social Cognitive and Affective Neuroscience* 11 (11). ISSN: 17495024. DOI: 10.1093/scan/nsw082.
- Hogg, Michael A., Deborah J. Terry and Katherine M. White (1995). 'A Tale of Two Theories: A Critical Comparison of Identity Theory with Social Identity Theory'. In: *Social Psychology Quarterly* 58 (4). ISSN: 01902725. DOI: 10.2307/2787127.
- Howland, John G. et al. (2022). 'The rodent medial prefrontal cortex and associated circuits in orchestrating adaptive behavior under variable demands'. In: *Neuroscience and Biobehavioral Reviews* 135. ISSN: 18737528. DOI: 10.1016/j.neubiorev.2022.104569.
- Hoyle, Rick H. et al. (2002). 'Reliability and validity of a brief measure of sensation seeking'. In: *Personality and Individual Differences* 32 (3). ISSN: 01918869. DOI: 10.1016/S0191-8869(01)00032-0.
- Hsu, Ming et al. (2009). 'Neural response to reward anticipation under risk is nonlinear in probabilities'. In: *Journal of Neuroscience* 29 (7). ISSN: 02706474. DOI: 10.1523/JNEUROSCI.5296-08.2009.
- Huettel, Scott A., Allen W. Song and Gregory McCarthy (2005). 'Decisions under uncertainty: Probabilistic context influences activation of prefrontal and parietal cortices'. In: *Journal of Neuroscience* 25 (13). ISSN: 02706474. DOI: 10.1523/JNEUROSCI.5070-04.2005.
- Huettel, Scott A. et al. (2006). 'Neural signatures of economic preferences for risk and ambiguity'. In: *Neuron* 49 (5). ISSN: 08966273. DOI: 10.1016/j.neuron.2006.01.024.
- Iacoboni, Marco (2009). 'Imitation, empathy, and mirror neurons'. In: *Annual Review of Psychology* 60. ISSN: 00664308. DOI: 10.1146/annurev.psych.60.110707.163604.
- Jacoby, Jacob and Leon B. Kaplan (1972). 'The Components of Perceived Risk'. In: *Sv* (January 1972). ISSN: 0021-9010.
- Jaeger, Bastian et al. (2019). 'Explaining the persistent influence of facial cues in social decision-making'. In: *Journal of Experimental Psychology: General* 148 (6). ISSN: 00963445. DOI: 10.1037/xge0000591.

- Jamali, Mohsen et al. (2021). 'Single-neuronal predictions of others' beliefs in humans'. In: *Nature* 591 (7851). ISSN: 14764687. DOI: 10.1038/s41586-021-03184-0.
- Jannati, Sima, Alok Kumar and Justin Wolfers (2016). 'In-Group Bias in Financial Markets'. In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.2884218.
- Jenni, Nicole L., Joshua D. Larkin and Stan B. Floresco (2017). 'Prefrontal dopamine D1 and D2 receptors regulate dissociable aspects of decision making via distinct ventral striatal and amygdalar circuits'. In: *Journal of Neuroscience* 37 (26). ISSN: 15292401. DOI: 10.1523/JNEUROSCI.0030-17.2017.
- Joseph, Elizabeth D. and Don C. Zhang (2021). 'Personality Profile of Risk-Takers: An Examination of the Big Five Facets'. In: *Journal of Individual Differences* 42 (4). ISSN: 21512299. DOI: 10.1027/1614-0001/a000346.
- Joutsa, Juho et al. (2012). 'Mesolimbic dopamine release is linked to symptom severity in pathological gambling'. In: *NeuroImage* 60 (4). ISSN: 10538119. DOI: 10.1016/j.neuroimage.2012.02.006.
- Kacelnik, Alex and Melissa Bateson (1996). 'Risky theories - The effects of variance on foraging decisions'. In: *American Zoologist* 36 (4). ISSN: 00031569. DOI: 10.1093/icb/36.4.402.
- Kahneman, Daniel (2012). *Thinking, fast and slow / Daniel Kahneman*.
- Kahneman, Daniel and Amos Tversky (1979). 'Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-291.' In: *Econometrica* 47 (2). ISSN: 00129682.
- (2018). *Prospect theory: An analysis of decision under risk*. Vol. 1. DOI: 10.2307/1914185.
- Kaufman, Alexandra Mansell, Tristan Geiller and Attila Losonczy (2020). 'A Role for the Locus Coeruleus in Hippocampal CA1 Place Cell Reorganization during Spatial Reward Learning'. In: *Neuron* 105 (6). ISSN: 10974199. DOI: 10.1016/j.neuron.2019.12.029.
- Király, Bálint and Balázs Hangya (2022). 'Navigating the Statistical Minefield of Model Selection and Clustering in Neuroscience'. In: *Eneuro* 9.4.

- Kirschner, Matthias et al. (2020). 'From apathy to addiction: Insights from neurology and psychiatry'. In: *Progress in Neuro-Psychopharmacology and Biological Psychiatry* 101. ISSN: 18784216. DOI: 10.1016/j.pnpbp.2020.109926.
- Kiverstein, Julian and Mark Miller (2015). 'The embodied brain: Towards a radical embodied cognitive neuroscience'. In: *Frontiers in Human Neuroscience* 9 (MAY). ISSN: 16625161. DOI: 10.3389/fnhum.2015.00237.
- Kleef, Gerben A. Van, Carsten K.W. De Dreu and Antony S.R. Manstead (2004). 'The interpersonal effects of emotions in negotiations: A motivated information processing approach'. In: *Journal of Personality and Social Psychology* 87 (4). ISSN: 00223514. DOI: 10.1037/0022-3514.87.4.510.
- Kobayashi, Shunsuke, Ofelia Pinto De Carvalho and Wolfram Schultz (2010). 'Adaptation of reward sensitivity in orbitofrontal neurons'. In: *Journal of Neuroscience* 30 (2). ISSN: 02706474. DOI: 10.1523/JNEUROSCI.4009-09.2010.
- Lak, Armin, William R Stauffer and Wolfram Schultz (2014). 'Dopamine prediction error responses integrate subjective value from different reward dimensions'. In: *Proceedings of the National Academy of Sciences* 111 (6), pp. 2343–2348. ISSN: 0027-8424.
- Lammel, Stephan, Byung Kook Lim and Robert C. Malenka (2014). 'Reward and aversion in a heterogeneous midbrain dopamine system'. In: *Neuropharmacology* 76 (PART B). ISSN: 00283908. DOI: 10.1016/j.neuropharm.2013.03.019.
- Larsen, Tobias and John P. O'Doherty (2014). 'Uncovering the spatio-temporal dynamics of value-based decision-making in the human brain: A combined fmri-EEG study'. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 369 (1655). ISSN: 14712970. DOI: 10.1098/rstb.2013.0473.
- Lau, Brian and Paul W. Glimcher (2008). 'Value Representations in the Primate Striatum during Matching Behavior'. In: *Neuron* 58 (3). ISSN: 08966273. DOI: 10.1016/j.neuron.2008.02.021.
- Lauffs, Marc M et al. (2020). 'Risk prediction error signaling: A two-component response?' In: *NeuroImage* 214, p. 116766. ISSN: 1053-8119.

- Leuchs, Laura, Max Schneider and Victor I. Spoormaker (2019). 'Measuring the conditioned response: A comparison of pupillometry, skin conductance, and startle electromyography'. In: *Psychophysiology* 56 (1). ISSN: 14698986. DOI: 10.1111/psyp.13283.
- Lewis, Melissa A. et al. (2008). 'Fitting In and Feeling Fine: Conformity and Coping Motives as Mediators of the Relationship Between Social Anxiety and Problematic Drinking'. In: *Psychology of Addictive Behaviors* 22 (1). ISSN: 0893164X. DOI: 10.1037/0893-164X.22.1.58.
- Lieberman, Matthew D. (2007). 'Social cognitive neuroscience: A review of core processes'. In: *Annual Review of Psychology* 58. ISSN: 00664308. DOI: 10.1146/annurev.psych.58.110405.085654.
- Littler, Dale and Demetris Melanthiou (2006). 'Consumer perceptions of risk and uncertainty and the implications for behaviour towards innovative retail services: The case of Internet Banking'. In: *Journal of Retailing and Consumer Services* 13 (6). ISSN: 09696989. DOI: 10.1016/j.jretconser.2006.02.006.
- Liu, Xiang and John W Polak (2007). 'Nonlinearity and specification of attitudes toward risk in discrete choice models'. In: *Transportation research record* 2014.1, pp. 27–31.
- Lo, Chung Chuan and Xiao Jing Wang (2006). 'Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks'. In: *Nature Neuroscience* 9 (7). ISSN: 10976256. DOI: 10.1038/nn1722.
- Lo, Steson and Sally Andrews (2015). 'To transform or not to transform: Using generalized linear mixed models to analyse reaction time data'. In: *Frontiers in psychology* 6, p. 1171.
- Lockwood, Patricia L., Matthew A.J. Apps and Steve W.C. Chang (2020). 'Is There a 'Social' Brain? Implementations and Algorithms'. In: *Trends in Cognitive Sciences* 24 (10). ISSN: 1879307X. DOI: 10.1016/j.tics.2020.06.011.
- Loewenstein, George F et al. (2001). 'Risk as feelings.' In: *Psychological bulletin* 127.2, p. 267.

- Loued-Khenissi, Leyla et al. (2020). 'Anterior insula reflects surprise in value-based decision-making and perception'. In: *NeuroImage* 210. ISSN: 10959572. DOI: 10.1016/j.neuroimage.2020.116549.
- Lowet, Adam S. et al. (2020). 'Distributional Reinforcement Learning in the Brain'. In: *Trends in Neurosciences* 43 (12). ISSN: 1878108X. DOI: 10.1016/j.tins.2020.09.004.
- Ludemann, Pamela M. and Charles A. Nelson (1988). 'Categorical Representation of Facial Expressions by 7-Month-Old Infants'. In: *Developmental Psychology* 24 (4). ISSN: 00121649. DOI: 10.1037/0012-1649.24.4.492.
- Maass, Wolfgang (1997). 'Fast sigmoidal networks via spiking neurons'. In: *Neural Computation* 9 (2). ISSN: 08997667. DOI: 10.1162/neco.1997.9.2.279.
- Mackie, Diane M., Leila T. Worth and Arlene G. Asuncion (1990). 'Processing of Persuasive In-Group Messages'. In: *Journal of Personality and Social Psychology* 58 (5). ISSN: 00223514. DOI: 10.1037/0022-3514.58.5.812.
- Mackintosh, N. J. (1975). 'A theory of attention: Variations in the associability of stimuli with reinforcement'. In: *Psychological Review* 82 (4). ISSN: 0033295X. DOI: 10.1037/h0076778.
- Macpherson, Tom and Takatoshi Hikida (2019). 'Role of basal ganglia neurocircuitry in the pathology of psychiatric disorders'. In: *Psychiatry and Clinical Neurosciences* 73 (6). ISSN: 14401819. DOI: 10.1111/pcn.12830.
- Makowski, Dominique et al. (2023). 'Automated Results Reporting as a Practical Tool to Improve Reproducibility and Methodological Best Practices Adoption'. In: *CRAN*. URL: <https://easystats.github.io/report/>.
- Markowitz, Harry (1952). 'The utility of wealth'. In: *Journal of political Economy* 60.2, pp. 151–158.
- Mata, Rui et al. (2018). 'Risk preference: A view from psychology'. In: *Journal of Economic Perspectives* 32 (2). ISSN: 08953309. DOI: 10.1257/jep.32.2.155.
- McClure, Samuel M., Nathaniel D. Daw and P. Read Montague (2003). 'A computational substrate for incentive salience'. In: *Trends in Neurosciences* 26 (8). ISSN: 01662236. DOI: 10.1016/S0166-2236(03)00177-2.

- Meertens, Ree M. and René Lion (2008). ‘Measuring an individual’s tendency to take risks: The risk propensity scale’. In: *Journal of Applied Social Psychology* 38 (6). ISSN: 00219029. DOI: 10.1111/j.1559-1816.2008.00357.x.
- Mikhael, John G. and Rafal Bogacz (2016). ‘Learning Reward Uncertainty in the Basal Ganglia’. In: *PLoS Computational Biology* 12 (9). ISSN: 15537358. DOI: 10.1371/journal.pcbi.1005062.
- Mink, Jonathan W. (1996). ‘The basal ganglia: Focused selection and inhibition of competing motor programs’. In: *Progress in Neurobiology* 50 (4). ISSN: 03010082. DOI: 10.1016/S0301-0082(96)00042-1.
- Mishra, Sandeep (2014). ‘Decision-Making Under Risk: Integrating Perspectives From Biology, Economics, and Psychology’. In: *Personality and Social Psychology Review* 18 (3). ISSN: 15327957. DOI: 10.1177/1088868314530517.
- Mishra, Sandeep and Martin L Lalumière (2011). ‘Individual differences in risk-propensity: Associations between personality and behavioral measures of risk’. In: *Personality and Individual Differences* 50.6, pp. 869–873.
- Moeller, Moritz, Sanjay Manohar and Rafal Bogacz (2022). ‘Uncertainty-guided learning with scaled prediction errors in the basal ganglia’. In: *PLoS Computational Biology* 18.5, e1009816.
- Moeller, Moritz et al. (2021). ‘An association between prediction errors and risk-seeking: Theory and behavioral evidence’. In: *PLoS Computational Biology* 17 (7). ISSN: 15537358. DOI: 10.1371/journal.pcbi.1009213.
- Mohr, Peter N.C. et al. (2010). ‘Neural foundations of risk-return trade-off in investment decisions’. In: *NeuroImage* 49 (3). ISSN: 10538119. DOI: 10.1016/j.neuroimage.2009.10.060.
- Molenberghs, Pascal et al. (2013). ‘Seeing is believing: Neural mechanisms of action-perception are biased by team membership’. In: *Human Brain Mapping* 34 (9). ISSN: 10659471. DOI: 10.1002/hbm.22044.
- Monosov, Ilya E. (Oct. 2020). ‘How Outcome Uncertainty Mediates Attention, Learning, and Decision-Making’. In: *Trends in Neurosciences* 43 (10), pp. 795–809. ISSN: 1878108X. DOI: 10.1016/J.TINS.2020.06.009.

- Montague, P. Read, Peter Dayan and Terrence J. Sejnowski (1996). 'A framework for mesencephalic dopamine systems based on predictive Hebbian learning'. In: *Journal of Neuroscience* 16 (5). ISSN: 02706474. DOI: 10.1523/jneurosci.16-05-01936.1996.
- Morris, Genela et al. (2006). 'Midbrain dopamine neurons encode decisions for future action'. In: *Nature Neuroscience* 9 (8). ISSN: 10976256. DOI: 10.1038/nn1743.
- Mulcahy, Garrett, Brady Atwood and Alexey Kuznetsov (2020). 'Basal ganglia role in learning rewarded actions and executing previously learned choices: Healthy and diseased states'. In: *PLoS ONE* 15 (2). ISSN: 19326203. DOI: 10.1371/journal.pone.0228081.
- Mullen, Brian, Birgit Bryant and James E. Driskell (1997). 'Presence of others and arousal: An integration'. In: *Group Dynamics* 1 (1). ISSN: 10892699. DOI: 10.1037/1089-2699.1.1.52.
- Murayama, Kou and Shinji Kitagami (2014). 'Consolidation power of extrinsic rewards: Reward cues enhance long-term memory for irrelevant past events'. In: *Journal of Experimental Psychology: General* 143 (1). ISSN: 00963445. DOI: 10.1037/a0031992.
- Murty, Vishnu P., Kevin S. LaBar and R. Alison Adcock (2012). 'Threat of punishment motivates memory encoding via amygdala, not midbrain, interactions with the medial temporal lobe'. In: *Journal of Neuroscience* 32 (26). ISSN: 02706474. DOI: 10.1523/JNEUROSCI.0094-12.2012.
- Myers, David G (2016). *SOCIAL PSYCHOLOGY, TWELFTH EDITION*. Vol. 66.
- Nakatani, Yoshihiro et al. (2009). 'Why the carrot is more effective than the stick: Different dynamics of punishment memory and reward memory and its possible biological basis'. In: *Neurobiology of Learning and Memory* 92 (3). ISSN: 10747427. DOI: 10.1016/j.nlm.2009.05.003.
- Nassar, Matthew R et al. (2010). 'An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment'. In: *Journal of Neuroscience* 30.37, pp. 12366–12378.
- NIH (Dec. 2022). <https://www.nimh.nih.gov/health/statistics/any-anxiety-disorder>. URL: <https://www.nimh.nih.gov/health/statistics/any-anxiety-disorder>.

- Niv, Yael, Michael O Duff and Peter Dayan (2005). ‘Dopamine, uncertainty and TD learning’. In: *Behavioral and brain Functions* 1 (1), pp. 1–9. ISSN: 1744-9081.
- Niv, Yael and P. Read Montague (2009). ‘Theoretical and empirical studies of learning’. In: *Neuroeconomics*. DOI: 10.1016/B978-0-12-374176-9.00022-1.
- Niv, Yael et al. (2012). ‘Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain’. In: *Journal of Neuroscience* 32 (2). ISSN: 02706474. DOI: 10.1523/JNEUROSCI.5498-10.2012.
- Nobanee, Haitham et al. (2021). ‘A bibliometric analysis of objective and subjective risk’. In: *Risks* 9.7, p. 128.
- Novak, Keisha D. and Dan Foti (2015). ‘Teasing apart the anticipatory and consummatory processing of monetary incentives: An event-related potential study of reward dynamics’. In: *Psychophysiology* 52 (11). ISSN: 14698986. DOI: 10.1111/psyp.12504.
- Oakes, Margaret and Robert Bor (2010). ‘The psychology of fear of flying (part I): A critical evaluation of current perspectives on the nature, prevalence and etiology of fear of flying’. In: *Travel medicine and infectious disease* 8.6, pp. 327–338.
- Ochsner, Kevin N and Matthew D Lieberman (2001). ‘The emergence of social cognitive neuroscience.’ In: *American psychologist* 56.9, p. 717.
- O’Doherty, John P. et al. (Apr. 2003). ‘Temporal Difference Models and Reward-Related Learning in the Human Brain’. In: *Neuron* 38 (2), pp. 329–337. ISSN: 0896-6273. DOI: 10.1016/S0896-6273(03)00169-7.
- Olschewski, Sebastian, Linan Diao and Jörg Rieskamp (2021). ‘Reinforcement learning about asset variability and correlation in repeated portfolio decisions’. In: *Journal of Behavioral and Experimental Finance* 32, p. 100559.
- O’Neill, Martin and Wolfram Schultz (2010a). ‘Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value’. In: *Neuron* 68 (4). ISSN: 08966273. DOI: 10.1016/j.neuron.2010.09.031.
- (2010b). ‘Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value’. In: *Neuron* 68.4, pp. 789–800.
- (2013). ‘Risk prediction error coding in orbitofrontal neurons’. In: *Journal of Neuroscience* 33 (40), pp. 15810–15814. ISSN: 0270-6474.

- Ott, Torben and Andreas Nieder (2019). ‘Dopamine and cognitive control in prefrontal cortex’. In: *Trends in cognitive sciences* 23.3, pp. 213–234.
- Ouden, Hanneke E.M. Den, Peter Kok and Floris P. de Lange (2012). ‘How prediction errors shape perception, attention, and motivation’. In: *Frontiers in Psychology* 3 (DEC). ISSN: 16641078. DOI: 10.3389/fpsyg.2012.00548.
- Ouyang, Guang et al. (2017). ‘Exploiting the intra-subject latency variability from single-trial event-related potentials in the P3 time range: A review and comparative evaluation of methods’. In: *Neuroscience and Biobehavioral Reviews* 75. ISSN: 18737528. DOI: 10.1016/j.neubiorev.2017.01.023.
- O’Neill, Martin and Wolfram Schultz (2015). ‘Economic risk coding by single neurons in the orbitofrontal cortex’. In: *Journal of Physiology-Paris* 109 (1-3), pp. 70–77. ISSN: 0928-4257.
- Padoa-Schioppa, Camillo (2009). ‘Range-adapting representation of economic value in the orbitofrontal cortex’. In: *Journal of Neuroscience* 29 (44). ISSN: 02706474. DOI: 10.1523/JNEUROSCI.3751-09.2009.
- Pagnoni, Giuseppe et al. (Jan. 2002). ‘Activity in human ventral striatum locked to errors of reward prediction’. In: *Nature Neuroscience* 2002 5:2 5 (2), pp. 97–98. ISSN: 1546-1726. DOI: 10.1038/nn802. URL: <https://www.nature.com/articles/nn802>.
- Park, Soyoung Q. et al. (2012). ‘Adaptive coding of reward prediction errors is gated by striatal coupling’. In: *Proceedings of the National Academy of Sciences of the United States of America* 109 (11). ISSN: 00278424. DOI: 10.1073/pnas.1119969109.
- Parks, Craig D. and Lawrence J. Sanna (2018). ‘Group Influence’. In: *Group Performance and Interaction*. DOI: 10.4324/9780429500091-3.
- Parra, Lucas C. et al. (2003). ‘Response error correction - A demonstration of improved human-machine performance using real-time EEG monitoring’. In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 11 (2). ISSN: 15344320. DOI: 10.1109/TNSRE.2003.814446.
- (Nov. 2005). ‘Recipes for the linear analysis of EEG’. In: *NeuroImage* 28 (2), pp. 326–341. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2005.05.032.

- Paulhus, DL (2002). 'Socially desirable responding: The evolution of a construct'. In: *The role of constructs in psychological and ...* (1958).
- Payzan-LeNestour, Elise et al. (2013). 'The Neural Representation of Unexpected Uncertainty during Value-Based Decision Making'. In: *Neuron* 79 (1). ISSN: 08966273. DOI: 10.1016/j.neuron.2013.04.037.
- Pearce, John M. and Geoffrey Hall (1980). 'A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli'. In: *Psychological Review* 87 (6). ISSN: 0033295X. DOI: 10.1037/0033-295X.87.6.532.
- Peirce, Jonathan et al. (2019). 'PsychoPy2: Experiments in behavior made easy'. In: *Behavior Research Methods* 51 (1), pp. 195–203. ISSN: 1554-3528. DOI: 10.3758/s13428-018-01193-y. URL: <https://doi.org/10.3758/s13428-018-01193-y>.
- Pessiglione, Mathias et al. (2006). 'Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans'. In: *Nature* 442 (7106). ISSN: 14764687. DOI: 10.1038/nature05051.
- Pfabigan, Daniela M. et al. (2014). 'P300 amplitude variation is related to ventral striatum BOLD response during gain and loss anticipation: An EEG and fMRI experiment'. In: *NeuroImage* 96. ISSN: 10959572. DOI: 10.1016/j.neuroimage.2014.03.077.
- Philiastides, Marios G., Guido Biele and Hauke R. Heekeren (2010a). 'A mechanistic account of value computation in the human brain'. In: *Proceedings of the National Academy of Sciences of the United States of America* 107 (20). ISSN: 00278424. DOI: 10.1073/pnas.1001732107.
- Philiastides, Marios G. and Paul Sajda (2006). 'Temporal characterization of the neural correlates of perceptual decision making in the human brain'. In: *Cerebral Cortex* 16 (4). ISSN: 10473211. DOI: 10.1093/cercor/bhi130.
- Philiastides, Marios G. et al. (2010b). 'Temporal dynamics of prediction error processing during reward-based decision making'. In: *NeuroImage* 53 (1). ISSN: 10538119. DOI: 10.1016/j.neuroimage.2010.05.052.

- Phillips, Paul E.M., Mark E. Walton and Thomas C. Jhou (2007). 'Calculating utility: Preclinical evidence for cost-benefit analysis by mesolimbic dopamine'. In: *Psychopharmacology* 191 (3). ISSN: 00333158. DOI: 10.1007/s00213-006-0626-6.
- Piray, Payam and Nathaniel D. Daw (2021). 'A model for learning based on the joint estimation of stochasticity and volatility'. In: *Nature Communications* 12 (1). ISSN: 20411723. DOI: 10.1038/s41467-021-26731-9.
- Platt, Michael L and Scott A Huettel (2008). 'Risky business: the neuroeconomics of decision making under uncertainty'. In: *Nature neuroscience* 11 (4), pp. 398–403. ISSN: 1546-1726.
- Plötner, Maria et al. (2015). 'The effects of collaboration and minimal-group membership on children's prosocial behavior, liking, affiliation, and trust'. In: *Journal of Experimental Child Psychology* 139. ISSN: 00220965. DOI: 10.1016/j.jecp.2015.05.008.
- Porcelli, Anthony J. and Mauricio R. Delgado (2017). 'Stress and decision making: effects on valuation, learning, and risk-taking'. In: *Current Opinion in Behavioral Sciences* 14. ISSN: 23521546. DOI: 10.1016/j.cobeha.2016.11.015.
- Porter, Stephen, Leanne ten Brinke and Chantal Gustaw (2010). 'Dangerous decisions: The impact of first impressions of trustworthiness on the evaluation of legal evidence and defendant culpability'. In: *Psychology, Crime and Law* 16 (6). ISSN: 1068316X. DOI: 10.1080/10683160902926141.
- Potenza, Marc N. (2008). 'The neurobiology of pathological gambling and drug addiction: An overview and new findings'. In: vol. 363. DOI: 10.1098/rstb.2008.0100.
- Powers, Katherine E. et al. (2022). 'Effects of peer observation on risky decision-making in adolescence: A meta-analytic review'. In: *Psychological Bulletin* 148 (11-12). ISSN: 0033-2909. DOI: 10.1037/bu10000382.
- Premack, David and Guy Woodruff (1978). 'Does the chimpanzee have a theory of mind?' In: *Behavioral and Brain Sciences* 1 (4). ISSN: 14691825. DOI: 10.1017/S0140525X00076512.

- Prentice-Dunn, Steven and Ronald W. Rogers (1982). 'Effects of public and private self-awareness on deindividuation and aggression'. In: *Journal of Personality and Social Psychology* 43 (3). ISSN: 00223514. DOI: 10.1037/0022-3514.43.3.503.
- Preuschoff, Kerstin and Peter Bossaerts (2007). 'Adding prediction risk to the theory of reward learning'. In: vol. 1104. DOI: 10.1196/annals.1390.005.
- Preuschoff, Kerstin, Peter Bossaerts and Steven R. Quartz (2006). 'Neural Differentiation of Expected Reward and Risk in Human Subcortical Structures'. In: *Neuron* 51 (3). ISSN: 08966273. DOI: 10.1016/j.neuron.2006.06.024.
- Preuschoff, Kerstin, Bernard Marius 't Hart and Wolfgang Einhäuser (2011). 'Pupil dilation signals surprise: Evidence for noradrenaline's role in decision making'. In: *Frontiers in neuroscience* 5, p. 115.
- Preuschoff, Kerstin, Steven R Quartz and Peter Bossaerts (2008). 'Human insula activation reflects risk prediction errors as well as risk'. In: *Journal of Neuroscience* 28 (11), pp. 2745–2752. ISSN: 0270-6474.
- Proske, Dirk and Dirk Proske (2008). 'Objective Risk Measures'. In: *Catalogue of Risks: Natural, Technical, Social and Health Risks*, pp. 267–327.
- Qi, Yue et al. (2022). 'Trustworthy faces make people more risk-tolerant: The effect of facial trustworthiness on risk decision-making under gain and loss conditions'. In: *PsyCh Journal* 11.1, pp. 43–50.
- Rangel, Antonio, Colin Camerer and P. Read Montague (2008). 'A framework for studying the neurobiology of value-based decision making'. In: *Nature Reviews Neuroscience* 9 (7). ISSN: 1471003X. DOI: 10.1038/nrn2357.
- Reimer, Jacob et al. (2016). 'Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex'. In: *Nature communications* 7 (1), pp. 1–7. ISSN: 2041-1723.
- Rejda, George E et al. (2005). 'Principles of risk management and insurance'. In.
- Rescorla, Robert A and Allan R Wagner (1972). 'A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement'. In: *Classical conditioning II: current research and theory*.

- Reyna, Valerie F. and Frank Farley (2006). 'Risk and rationality in adolescent decision making: Implications for theory, practice, and public policy'. In: *Psychological Science in the Public Interest, Supplement 7* (1). ISSN: 15291006. DOI: 10.1111/j.1529-1006.2006.00026.x.
- Reynolds, John N.J., Brian I. Hyland and Jeffery R. Wickens (2001). 'A cellular mechanism of reward-related learning'. In: *Nature* 413 (6851). ISSN: 00280836. DOI: 10.1038/35092560.
- Rezlescu, Constantin et al. (2012). 'Unfakeable facial configurations affect strategic choices in trust games with or without information about past behavior'. In: *PLoS ONE* 7 (3). ISSN: 19326203. DOI: 10.1371/journal.pone.0034293.
- Rilling, James K. and Alan G. Sanfey (2011). 'The neuroscience of social decision-making'. In: *Annual Review of Psychology* 62. ISSN: 00664308. DOI: 10.1146/annurev.psych.121208.131647.
- Robbins, T. W. and B. J. Everitt (2007). 'A role for mesencephalic dopamine in activation: Commentary on Berridge (2006)'. In: *Psychopharmacology* 191 (3). ISSN: 00333158. DOI: 10.1007/s00213-006-0528-7.
- Roesch, Matthew R. and Carl R. Olson (2004). 'Neuronal Activity Related to Reward Value and Motivation in Primate Frontal Cortex'. In: *Science* 304 (5668). ISSN: 00368075. DOI: 10.1126/science.1093223.
- Rogers-Carter, Morgan M and John P Christianson (2019). 'An insular view of the social decision-making network'. In: *Neuroscience & Biobehavioral Reviews* 103, pp. 119–132.
- Rolls, Edmund T., Ciara McCabe and Jerome Redoute (2008). 'Expected value, reward outcome, and temporal difference error representations in a probabilistic decision task'. In: *Cerebral Cortex* 18 (3). ISSN: 10473211. DOI: 10.1093/cercor/bhm097.
- Rouhani, Nina and Yael Niv (2021). 'Signed and unsigned reward prediction errors dynamically enhance learning and memory'. In: *eLife* 10. ISSN: 2050084X. DOI: 10.7554/eLife.61077.

- Rousselet, Guillaume A, Cyril R Pernet and Rand R Wilcox (2019). ‘A practical introduction to the bootstrap: a versatile method to make inferences by using data-driven simulations’. In: *PsyArXiv Preprints*.
- Ruff, Christian C and Ernst Fehr (2014). ‘The neurobiology of rewards and values in social decision making’. In: *Nature Reviews Neuroscience* 15 (8), pp. 549–562. ISSN: 1471-0048.
- Rushworth, Matthew F.S. and Timothy E.J. Behrens (2008). ‘Choice, uncertainty and value in prefrontal and cingulate cortex’. In: *Nature Neuroscience* 11 (4). ISSN: 10976256. DOI: 10.1038/nn2066.
- Safarzynska, Karolina (2018). ‘The Impact of Resource Uncertainty and Intergroup Conflict on Harvesting in the Common-Pool Resource Experiment’. In: *Environmental and Resource Economics* 71 (4). ISSN: 15731502. DOI: 10.1007/s10640-017-0193-9.
- Sajda, Paul, Marios G. Philiastides and Lucas C. Parra (2009). ‘Single-Trial Analysis of Neuroimaging Data: Inferring Neural Networks Underlying Perceptual Decision-Making in the Human Brain’. In: *IEEE Reviews in Biomedical Engineering* 2. ISSN: 19411189. DOI: 10.1109/RBME.2009.2034535.
- Salamone, J. D. et al. (2005). ‘Beyond the reward hypothesis: Alternative functions of nucleus accumbens dopamine’. In: *Current Opinion in Pharmacology* 5 (1). ISSN: 14714892. DOI: 10.1016/j.coph.2004.09.004.
- Salamone, John D. and Mercè Correa (2002). ‘Motivational views of reinforcement: Implications for understanding the behavioral functions of nucleus accumbens dopamine’. In: *Behavioural Brain Research* 137 (1-2). ISSN: 01664328. DOI: 10.1016/S0166-4328(02)00282-6.
- (2012). ‘The Mysterious Motivational Functions of Mesolimbic Dopamine’. In: *Neuron* 76 (3). ISSN: 08966273. DOI: 10.1016/j.neuron.2012.10.021.
- Salamone, John D. et al. (2022). ‘Complexities and paradoxes in understanding the role of dopamine in incentive motivation and instrumental action: Exertion of effort vs. anhedonia’. In: *Brain Research Bulletin* 182. ISSN: 18732747. DOI: 10.1016/j.brainresbull.2022.01.019.

- Samejima, Kazuyuki et al. (2005). 'Neuroscience: Representation of action-specific reward values in the striatum'. In: *Science* 310 (5752). ISSN: 00368075. DOI: 10.1126/science.1115270.
- Sanfey, Alan G (2007). 'Social decision-making: insights from game theory and neuroscience'. In: *Science* 318 (5850), pp. 598–602. ISSN: 0036-8075.
- Santangelo, Gabriella et al. (2013). 'Pathological gambling in Parkinson's disease. A comprehensive review'. In: *Parkinsonism and Related Disorders* 19 (7). ISSN: 13538020. DOI: 10.1016/j.parkreldis.2013.02.007.
- Sara, Susan J. (2009). 'The locus coeruleus and noradrenergic modulation of cognition'. In: *Nature Reviews Neuroscience* 10 (3). ISSN: 1471003X. DOI: 10.1038/nrn2573.
- Scharlemann, Jörn P.W. et al. (2001). 'The value of a smile: Game theory with a human face'. In: *Journal of Economic Psychology* 22 (5). ISSN: 01674870. DOI: 10.1016/S0167-4870(01)00059-9.
- Schlicht, Erik J. et al. (2010). 'Human wagering behavior depends on opponents' faces'. In: *PLoS ONE* 5 (7). ISSN: 19326203. DOI: 10.1371/journal.pone.0011663.
- Schonberg, Tom, Craig R. Fox and Russell A. Poldrack (2011). 'Mind the gap: Bridging economic and naturalistic risk-taking with cognitive neuroscience'. In: *Trends in Cognitive Sciences* 15 (1). ISSN: 13646613. DOI: 10.1016/j.tics.2010.10.002.
- Schonberg, Tom et al. (2012). 'Decreasing ventromedial prefrontal cortex activity during sequential risk-taking: An fMRI investigation of the balloon analog risk task'. In: *Frontiers in Neuroscience* (JUN). ISSN: 16624548. DOI: 10.3389/fnins.2012.00080.
- Schuck-Paim, Cynthia, Lorena Pompilio and Alex Kacelnik (2004). 'State-dependent decisions cause apparent violations of rationality in animal choice'. In: *PLoS Biology* 2 (12). ISSN: 15449173. DOI: 10.1371/journal.pbio.0020402.
- Schultz, W. (2017a). 'Electrophysiological correlates of reward processing in dopamine neurons'. In: *Decision Neuroscience: An Integrative Perspective*. DOI: 10.1016/B978-0-12-805308-9.00002-6.

- Schultz, Wolfram (1997). ‘The Phasic Reward Signal of Primate Dopamine Neurons’. In: *Advances in Pharmacology* 42 (C). ISSN: 10543589. DOI: 10.1016/S1054-3589(08)60841-8.
- (2010). ‘Dopamine signals for reward value and risk: Basic and recent data’. In: *Behavioral and Brain Functions* 6. ISSN: 17449081. DOI: 10.1186/1744-9081-6-24.
- (2015). ‘Neuronal reward and decision signals: from theories to data’. In: *Physiological reviews* 95 (3), pp. 853–951. ISSN: 0031-9333.
- (2016a). ‘Dopamine reward prediction error coding’. In: *Dialogues in Clinical Neuroscience* 18 (1). ISSN: 12948322. DOI: 10.31887/dcns.2016.18.1/wschultz.
- (2016b). ‘Dopamine reward prediction-error signalling: a two-component response’. In: *Nature reviews neuroscience* 17 (3), p. 183. ISSN: 1471-0048.
- (2017b). ‘Neuronal risk processing in human and monkey prefrontal cortex’. In: *The Prefrontal Cortex as an Executive, Emotional, and Social Brain*. DOI: 10.1007/978-4-431-56508-6_6.
- (2017c). ‘Reward prediction error’. In: *Current Biology* 27 (10), R369–R371. ISSN: 0960-9822.
- Schultz, Wolfram, Peter Dayan and P Read Montague (1997). ‘A neural substrate of prediction and reward’. In: *Science* 275 (5306), pp. 1593–1599. ISSN: 0036-8075.
- Schultz, Wolfram et al. (2021). ‘Smarter than humans: rationality reflected in primate neuronal reward signals’. In: *Current Opinion in Behavioral Sciences* 41. ISSN: 23521546. DOI: 10.1016/j.cobeha.2021.03.021.
- Sesack, Susan R. and Anthony A. Grace (2010). ‘Cortico-basal ganglia reward network: Microcircuitry’. In: *Neuropsychopharmacology* 35 (1). ISSN: 0893133X. DOI: 10.1038/npp.2009.93.
- Sims, Chris R. et al. (2013). ‘Melioration as rational choice: Sequential decision making in uncertain environments’. In: *Psychological Review* 120 (1). ISSN: 0033295X. DOI: 10.1037/a0030850.
- Singmann, Henrik et al. (2016). *afex: Analysis of Factorial Experiments*. R package version 0.16-1. URL: <https://CRAN.R-project.org/package=afex>.

- Soltani, Alireza and Alicia Izquierdo (2019). 'Adaptive learning under expected and unexpected uncertainty'. In: *Nature Reviews Neuroscience* 20.10, pp. 635–644.
- Somerville, Leah H. et al. (2013). 'The Medial Prefrontal Cortex and the Emergence of Self-Conscious Emotion in Adolescence'. In: *Psychological Science* 24 (8). ISSN: 14679280. DOI: 10.1177/0956797613475633.
- Somerville, Leah H. et al. (2019). 'Dissecting "Peer Presence" and "Decisions" to Deepen Understanding of Peer Influence on Adolescent Risky Choice'. In: *Child Development* 90 (6). ISSN: 14678624. DOI: 10.1111/cdev.13081.
- Stauffer, William R., Armin Lak and Wolfram Schultz (2014). 'Dopamine reward prediction error responses reflect marginal utility'. In: *Current Biology* 24 (21). ISSN: 18790445. DOI: 10.1016/j.cub.2014.08.064.
- Stauffer, William R. et al. (2015). 'Economic choices reveal probability distortion in macaque monkeys'. In: *Journal of Neuroscience* 35 (7). ISSN: 15292401. DOI: 10.1523/JNEUROSCI.3653-14.2015.
- Steele, Claude M (1997). 'A threat in the air: How stereotypes shape intellectual identity and performance.' In: *American psychologist* 52.6, p. 613.
- Stein, Murray B. et al. (2007). 'Increased amygdala and insula activation during emotion processing in anxiety-prone subjects'. In: *American Journal of Psychiatry* 164 (2). ISSN: 0002953X. DOI: 10.1176/ajp.2007.164.2.318.
- Steinbach, M. C. (2001). 'Markowitz revisited: Mean-variance models in financial portfolio analysis'. In: *SIAM Review* 43 (1). ISSN: 00361445. DOI: 10.1137/S0036144500376650.
- Steinberg, Laurence (2008). 'A social neuroscience perspective on adolescent risk-taking'. In: *Developmental Review* 28 (1). ISSN: 02732297. DOI: 10.1016/j.dr.2007.08.002.
- Stojanović, Nikola M et al. (2020). 'Reliability and validity of the Spielberger's State-Trait Anxiety Inventory (STAI) in Serbian university student and psychiatric non-psychotic outpatient populations'. In: *Acta facultatis medicae Naissensis* 37.2, pp. 149–159.

- Stone, Eric R. and Liz Allgaier (2008). ‘A social values analysis of self-other differences in decision making involving risk’. In: *Basic and Applied Social Psychology* 30 (2). ISSN: 01973533. DOI: 10.1080/01973530802208832.
- Sussman, Abigail B., Kristina Petkova and Alexander Todorov (2013). ‘Competence ratings in US predict presidential election outcomes in Bulgaria’. In: *Journal of Experimental Social Psychology* 49 (4). ISSN: 00221031. DOI: 10.1016/j.jesp.2013.02.003.
- Sutton, Richard S and Andrew G Barto (2018). *Reinforcement learning: An introduction*. MIT press. ISBN: 0262352702.
- Suzuki, Shinsuke et al. (2012). ‘Learning to Simulate Others’ Decisions’. In: *Neuron* 74 (6). ISSN: 08966273. DOI: 10.1016/j.neuron.2012.04.030.
- Suzuki, Shinsuke et al. (2016). ‘Behavioral contagion during learning about another agent’s risk-preferences acts on the neural representation of decision-risk’. In: *Proceedings of the National Academy of Sciences of the United States of America* 113 (14). ISSN: 10916490. DOI: 10.1073/pnas.1600092113.
- Taghizadeh, Bahareh et al. (2020). ‘Reward uncertainty asymmetrically affects information transmission within the monkey fronto-parietal network’. In: *Communications Biology* 3 (1). ISSN: 23993642. DOI: 10.1038/s42003-020-01320-6.
- Tai, Lung Hao et al. (2012). ‘Transient stimulation of distinct subpopulations of striatal neurons mimics changes in action value’. In: *Nature Neuroscience* 15 (9). ISSN: 10976256. DOI: 10.1038/nn.3188.
- Takahashi, Hidehiko et al. (2009). ‘When your gain is my pain and your pain is my gain: Neural correlates of envy and schadenfreude’. In: *Science* 323 (5916). ISSN: 00368075. DOI: 10.1126/science.1165604.
- Team, R Core (2021). *R: A language and environment for statistical computing*.
- Terasawa, Yuri et al. (2021). ‘Effects of insular resection on interactions between cardiac interoception and emotion recognition’. In: *Cortex* 137. ISSN: 19738102. DOI: 10.1016/j.cortex.2021.01.011.

- Tobler, Philippe N et al. (2007). 'Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems'. In: *Journal of neurophysiology* 97 (2), pp. 1621–1632. ISSN: 0022-3077.
- Tobler, Philippe N. et al. (2009). 'Risk-dependent reward value signal in human prefrontal cortex'. In: *Proceedings of the National Academy of Sciences of the United States of America* 106 (17). ISSN: 00278424. DOI: 10.1073/pnas.0809599106.
- Todorov, Alexander, Manish Pakrashi and Nikolaas N. Oosterhof (2009). 'Evaluating faces on trustworthiness after minimal time exposure'. In: *Social Cognition* 27 (6). ISSN: 0278016X. DOI: 10.1521/soco.2009.27.6.813.
- Tom, Sabrina M et al. (2007). 'The neural basis of loss aversion in decision-making under risk'. In: *Science* 315.5811, pp. 515–518.
- Trepel, Christopher, Craig R. Fox and Russell A. Poldrack (2005). 'Prospect theory on the brain? Toward a cognitive neuroscience of decision under risk'. In: vol. 23. DOI: 10.1016/j.cogbrainres.2005.01.016.
- Vavra, Peter, Luke J Chang and Alan G Sanfey (2018). 'Expectations in the Ultimatum Game: distinct effects of mean and variance of expected offers'. In: *Frontiers in psychology* 9, p. 992. ISSN: 1664-1078.
- Voci, Alberto (June 2006). 'The link between identification and in-group favouritism: Effects of threat to social identity and trust-related emotions'. In: *British Journal of Social Psychology* 45 (2), pp. 265–284. ISSN: 2044-8309. DOI: 10.1348/014466605X52245.
- Volkow, Nora D., Roy A. Wise and Ruben Baler (2017). 'The dopamine motive system: Implications for drug and food addiction'. In: *Nature Reviews Neuroscience* 18 (12). ISSN: 14710048. DOI: 10.1038/nrn.2017.130.
- Wagatsuma, Akiko et al. (2017). 'Locus coeruleus input to hippocampal CA3 drives single-trial learning of a novel context'. In: *Proceedings of the National Academy of Sciences of the United States of America* 115 (2). ISSN: 10916490. DOI: 10.1073/pnas.1714082115.
- Wake, Stephanie J. and Keise Izuma (2017). 'A common neural code for social and monetary rewards in the human striatum'. In: *Social Cognitive and Affective Neuroscience* 12 (10). ISSN: 17495024. DOI: 10.1093/scan/nsx092.

- Wassum, Kate M. et al. (2013). 'Phasic mesolimbic dopamine release tracks reward seeking during expression of pavlovian-to-instrumental transfer'. In: *Biological Psychiatry* 73 (8). ISSN: 00063223. DOI: 10.1016/j.biopsych.2012.12.005.
- Watabe-Uchida, Mitsuko, Neir Eshel and Naoshige Uchida (2017). 'Neural Circuitry of Reward Prediction Error'. In: *Annual Review of Neuroscience* 40. ISSN: 15454126. DOI: 10.1146/annurev-neuro-072116-031109.
- Weigard, Alexander et al. (2014). 'Effects of anonymous peer observation on adolescents' preference for immediate rewards'. In: *Developmental science* 17.1, pp. 71–78.
- Williams, Mark A. et al. (2004). 'Amygdala Responses to Fearful and Happy Facial Expressions under Conditions of Binocular Suppression'. In: *Journal of Neuroscience* 24 (12). ISSN: 02706474. DOI: 10.1523/JNEUROSCI.4977-03.2004.
- Wolf, Laura K et al. (2015). 'The audience effect in adolescence depends on who's looking over your shoulder'. In: *Journal of adolescence* 43, pp. 5–14.
- Wolfschlag, Mirjam and Anders Håkansson (2023). 'Drug-Induced Gambling Disorder: Epidemiology, Neurobiology, and Management'. In: *Pharmaceutical Medicine* 37 (1). ISSN: 11791993. DOI: 10.1007/s40290-022-00453-9.
- Wood, Wendy and Timothy Hayes (2012). 'Social Influence on consumer decisions: Motives, modes, and consequences'. In: *Journal of Consumer Psychology* 22 (3). ISSN: 10577408. DOI: 10.1016/j.jcps.2012.05.003.
- Wulff, Dirk U., Max Mergenthaler-Canseco and Ralph Hertwig (2018). 'A meta-analytic review of two modes of learning and the description-experience gap'. In: *Psychological Bulletin* 144 (2). ISSN: 00332909. DOI: 10.1037/bu10000115.
- Wunderlich, Klaus, Antonio Range and John P. O'Doherty (2009). 'Neural computations underlying action-based decision making in the human brain'. In: *Proceedings of the National Academy of Sciences of the United States of America* 106 (40). ISSN: 00278424. DOI: 10.1073/pnas.0901077106.
- Xiang, Ting, Terry Lohrenz and P. Read Montague (2013). 'Computational substrates of norms and their violations during social exchange'. In: *Journal of Neuroscience* 33 (3). ISSN: 02706474. DOI: 10.1523/JNEUROSCI.1642-12.2013.

- Yu, Hui, Oliver G.B. Garrod and Philippe G. Schyns (2012). 'Perception-driven facial expression synthesis'. In: vol. 36. DOI: 10.1016/j.cag.2011.12.002.
- Zhan, Jiayu et al. (2019). 'Modelling face memory reveals task-generalizable representations'. In: *Nature Human Behaviour* 3 (8). ISSN: 23973374. DOI: 10.1038/s41562-019-0625-3.
- Zhang, Jun Fang et al. (2021). 'Impulse Control Disorders in Parkinson's Disease: Epidemiology, Pathogenesis and Therapeutic Strategies'. In: *Frontiers in Psychiatry* 12. ISSN: 16640640. DOI: 10.3389/fpsyt.2021.635494.
- Zhang, Yuanyuan et al. (2017). 'Temporal dynamics of reward anticipation in the human brain'. In: *Biological Psychology* 128. ISSN: 18736246. DOI: 10.1016/j.biopsycho.2017.07.011.
- Zhao, Yingnan et al. (2020). 'Obtaining accurate estimated action values in categorical distributional reinforcement learning'. In: *Knowledge-Based Systems* 194. ISSN: 09507051. DOI: 10.1016/j.knosys.2020.105511.