

**Individual differences in value-based decision-making as
predictors of substance dependence**



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Laura Rai

Supervisor: Robert Whelan, PhD

A thesis submitted to the School of Psychology, Trinity College Dublin
In fulfilment of the requirement of the Degree of Doctor of Philosophy

2021

Table of Contents

Thesis Summary	vi
Publications arising from the thesis:	viii
Conference papers and peer-reviewed abstracts:.....	viii
Acknowledgements.....	ix
Acronyms & Abbreviations	x
Chapter 1: Literature Review	1
1. 1. Perspectives on substance dependence	2
1. 2. Neurobiology of Substance Dependence.....	7
1. 3. Substance dependence and decision-making.....	8
1. 3. 1. Alcohol.....	10
1. 3. 2. Opioids	10
1. 3. 3. Nicotine	11
1. 4. Electrophysiology	14
1. 5. Probabilistic Selection Task.....	18
1. 6. Computational Psychiatry	22
1. 6. 1. Reinforcement Learning.....	24
1. 6. 2. Computational Psychiatry & Addiction	28
1. 7. Drift-diffusion Model.....	31
1. 7. 1. Application of DDM	33
1. 7. 2. Combined RL and DDM models.....	34
1. 8. Bridging behavioural and neural data.....	34
1. 9. Machine Learning.....	35
1. 10. Unresolved Issues & Current Research	38
Chapter 2: Methodology	41
2. 1. Experimental Task	42
2. 1. 1. Probabilistic Selection Task (Presentation)	44
2. 1. 2. Probabilistic Selection Task (Inquisit).....	44
2. 2 Barratt Impulsivity Questionnaire	45
2. 3. Fagerstrom Test for Nicotine Dependence	46
2. 4. Computational Modelling of PST	46
2. 4. 1. Hierarchical Drift Diffusion Model	46
2. 4. 2. Statistical Approach.....	48
2. 5. Machine Learning.....	49
Chapter 3: Value-based decision-making as a predictor of hazardous alcohol-use	52

3. 1. Introduction	53
3. 2. Methods	55
3. 3. Results	59
3. 3. 1. Behavioral Results	59
3. 3. 2. Computational Models.....	62
3. 3. 3. Machine Learning.....	67
3. 3. 4. Correlations.....	77
3. 4. Discussion.....	85
Chapter 4: Reward and punishment processing in former opioid addiction	91
4.1. Introduction	92
4. 2. Methods	95
4. 2. 1. Sample	95
4. 2. 2. Materials	95
4. 2. 2. 1. Probabilistic Selection Task	95
4. 2. 2. 2. Barratt Impulsiveness Scale (BIS)	95
4. 2. 2. 3. Short- Impulsive Behavior Scale (UPPS-P).....	95
4. 2. 2. 4. Quality Control	95
4. 2. 2. 5. Computational Modelling	95
4. 2. 2. 6. Machine Learning.....	96
4. 3. Results	96
4. 3. 1. Behavioural Results	96
4. 3. 2. Computational Modelling Results.....	99
4. 3. 3. Posterior Predictive Checks.....	101
4. 3. 4. Machine Learning.....	103
4. 3. 5. Correlations with Length of Abstinence.....	113
4. 3. 6. Correlations between ML Features.....	114
4. 4. Discussion.....	116
Chapter 5: Reward and punishment learning predictors of smoking status	120
5. 1. Introduction	121
5. 2 Study 1.....	121
5. 2. 1. Methods	121
5. 2. 2. Results	123
5. 2. 3. Study 1 Discussion.....	126
5.3. Study 2.....	128
5.3.1. Methods	128
5. 3. 1. 1. Participants.....	128

5. 3. 1. 2. Materials.....	128
5. 3. 1. 2. 1. Machine Learning - Group Classification.....	129
5. 3. 2. Results	129
5. 3. 2. 1. Behavioral Results	129
5. 3. 2. 2. Modelling of PST Training Phase	134
5. 3. 2. 3. Modelling of PST Test Phase	135
5. 3. 2. 4. Group Classification: Current Smokers vs. Non-smokers.....	138
5. 3. 2. 5. Multinomial Group Classification – Current Smokers, Vapers, Ex-smokers, Non-smokers	143
5. 3. 2. 6. Correlations between ML Features	158
5.3.3. Study 2 - Discussion.....	160
5. 4. General Conclusion	162
Chapter 6: Longitudinal changes in value-based decision-making as predictors of nicotine relapse.....	164
6. 1. Introduction	165
6. 2. Methods	167
6. 2. 1. Participants.....	167
6. 2. 2. Materials	167
6. 2. 2. 1. Probabilistic Selection Task	167
6. 2. 2. 2. Barratt Impulsivity Scale.....	167
6. 2. 2. 3. Fagerstrom Test for Nicotine Dependence (FTND)	167
6. 2. 2. 4. Hooked on Nicotine Checklist (HONC)	168
6. 2. 2. 5. Short Form Smoking Consequences Questionnaire (S-SCQ).....	168
6. 2. 2. 6. EEG	168
6. 2. 3. Procedure	169
Behavioural data analysis.....	170
Drift-diffusion Models.....	170
6. 3. Results	170
6. 3. 1. Behavioural analysis	170
6. 3. 2. ERP analysis	172
Whole Group.....	173
ERP Group Comparisons	173
PST Test Phase ERPs.....	179
6. 3. 3. Computational Modelling Results.....	180
6. 3. 3. 1. Behavioural vs. Neural data RLHDDM Models.....	180
6. 3. 3. 2. Drift-diffusion Computational Models of PST	181

6. 3. 4. Cox Regression	182
6. 3. 5. Correlations between ML features	186
6. 4. Discussion.....	189
Chapter 7: General Discussion	192
7. 1. Summary of Results	194
7. 2. Relationship with existing literature	198
7. 3. Methodological Strengths & Limitations	209
7. 4. Conclusion	211
Appendix A: Chapter 3 Supplemental Results.....	213
A. 1. Correlations	213
A. 2. Non-hierarchical models.....	220
A. 3. Between-groups computational models	220
RL-HDDM.....	220
Between-groups HDDM	222
A. 4. Machine Learning Models with Residual Features.....	224
Appendix B: Chapter 4 Supplemental Results.....	227
B. 1. Correlations.....	227
B. 2. Comparison of HDDM Drift-rate Conditions.....	232
B. 3. Non-hierarchical RLDDM.....	232
Appendix C: Chapter 5 Supplementary Results.....	234
C. 1. Non-hierarchical models.....	234
Appendix D: Chapter 6 Supplementary Analysis.....	235
D. 1. Topoplots of EEG data	235
D. 1. Comparison of Computational Models.....	240
D. 2. Within-subjects of abstinence on boundary separation	241
D. 1. Supplementary Analysis of Event-related Potential-Reward Prediction Error Modulation	242
D. 1. 2. Reinforcement Learning	242
D. 1. 3. Single trial analysis.....	244
Whole Group.....	245
Single-trial Modulation Group Comparisons	247
D. 3. 3. Single-trial Cox Regression.....	250
D. 2. Correlation between PST at Study Time-points	250
Appendix E: Chapter 6 – Smoking Cessation Study Eligibility Criteria	251

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

X

Laura Rai

Signed:

Thesis Summary

Substance dependence is a leading global health concern. Alcohol use is the primary risk-factor for deaths and disability-adjusted life years among those between the ages of 15-49 (Griswold et al., 2018). Opioid dependence accounts for the majority of drug use disorders worldwide (James et al., 2018), while tobacco-use alone results in over one in ten deaths per year (Reitsma et al., 2017). The chronically relapsing nature of substance-dependence renders challenges for successful abstinence, with some studies reporting more than two thirds of individuals relapse within months of treatment initiation (Moeller & Paulus, 2018; Sinha, 2011). A key characteristic of substance-dependence is the persistent and continued use of substances despite their negative consequences. Therefore, understanding the psychological and neurobiological mechanisms that give rise to such maladaptive decision-making is a key goal for the field of addiction science. Converging evidence suggests that aberrant value-based decision-making (e.g., reward processing) characterises substance-dependent individuals from healthy controls, and may predict future use. However, it remains unclear whether different substance types, misuse patterns, and treatment interventions differentially affect decision-making impairments.

This thesis investigated value-based decision-making in various substance-dependence phenotypes. Specifically, computational models of decision-making (reinforcement learning and drift-diffusion models; RLDDMs) were fit to choice and reaction time data from a popular reward learning task known to index fluctuations dopaminergic functioning and show sensitivity to various clinical disorders (Frank et al., 2004). Machine learning (ML) methods were leveraged to investigate if parameters derived from the computational models could successfully predict substance-dependence. Four ML models were compared: (i) a Summary model with mean choice accuracy from the reward learning task, (ii) a Computational model with parameter estimates from RLDDMs, (iii) a Personality model with self-reported impulsivity, and (iv) a Combined model with features from (ii) and (iii). Additionally, each individual model was compared with a 'null' model including demographic features.

The literature on value-based decision-making and substance dependence was reviewed in Chapter 1. Chapter 2 described the general methods. Chapter 3 sought to predict hazardous alcohol-use risk (N=115). Chapter 4 sought to predict length of abstinence in a sample formerly dependent on heroin currently in methadone maintenance treatment (N=81). Chapter 5 sought to classify individuals based on their smoking group category (non-smokers, current smokers, ex-smokers, and vapers; N = 173). Chapter 6 aimed to assess behavioural and electrophysiological longitudinal changes in value-based decision-making during a smoking quit attempt (N = 112).

The results revealed reductions in response caution (indexed by the boundary separation parameter in drift-diffusion models) significantly predicted higher alcohol misuse risk, shorter lengths of opioid abstinence, and smoker versus non-smoker group membership. These findings suggest that response caution may be a task-general marker of substance-dependence that is sensitive to length of methadone treatment. Efficiency of evidence accumulation (i.e., the process of accumulating evidence for one option relative to another; drift-rate) was also a significant predictor across studies, however did not show a clear directional relationship with substance dependence, and was influenced by conflict in reward probabilities. For example, ex-smokers were classified by reduced evidence accumulation for high conflict trials with stimuli associated with negative feedback, and increased evidence accumulation for high conflict trials with stimuli associated with positive feedback.

Computational models performed similarly with Personality and Summary models, and outperformed demographic models overall. This indicates that self-reported impulsivity and mean choices in the decision-making task were as predictive compared with computational models fit to trial-by-trial choice and RT data. Overall, these findings highlight the utility of RLDDMs to investigate clinically relevant features of instrumental learning and decision-making, and identify features of value-based decision-making (i.e. evidence accumulation and decision threshold) that are predictive of substance-dependence.

Publications arising from the thesis:

Rai L. A, O'Halloran L, Jollans L, Vahey N, O'Brolchain C, Whelan R. Individual differences in learning from probabilistic reward and punishment predicts smoking status. (2019). *Addictive Behaviors*, 1;88:73–6. doi: 10.1016/j.addbeh.2018.08.019

Conference papers and peer-reviewed abstracts:

Rai, L. & Whelan, R. (2020). Reinforcement learning signals in value-based decision-making predict alcohol use. Presented at Neuromatch 2.0. Poster available at: <https://osf.io/vwm6c/>

Rai, L., Cao, Z., Ferdinand-Lespine, L., Rueda-Delgado, L., Vahey, N., Ruddy, K., (...), & Whelan, R. (2019). Examining neural and computational components of probabilistic decision-making in a longitudinal smoking cessation study. Presented at the 6th Model Based Neuroscience Summer School, Amsterdam, Netherlands. Poster available at: <https://osf.io/euxtr/>

Rai, L., Boyle, R., Tarnanas, I., & Whelan, R. (2019). Evaluation of a computerized cognitive screening tool (ADPS, Altoida) in a sample of community-dwelling older adults. 11th International Dementia Conference, Croke Park, Dublin.

Rai, L., & Whelan, R. (2019). Identifying neural and computational markers of age-related changes in reinforcement learning. Oral and poster presentation at the British Neuroscience Association Conference, Dublin.

Rai, L., & Whelan, R., (2018). Individual differences in learning from probabilistic reward and punishment predicts smoking status. Poster presented at the Young Neuroscience Ireland Symposium, University College Dublin.

Acknowledgements

This thesis would not have been possible without the support of many people and institutions over the past few years. Firstly, I would like to sincerely thank my research supervisor, Dr. Robert Whelan, for his academic guidance, wisdom, and patience throughout the whole PhD process. I feel extremely fortunate to have learned so much during my time in his lab and am grateful for the opportunity. I would also like to thank Dr. Ioannis Tarnanas and Altoida Inc. for their academic and funding support during my PhD term, and the Irish Research Council for supporting this research.

Thank you to Dr. Jo-Hanna Ivers and her research team in the School of Medicine, Trinity College Dublin, and my thesis appraisal committee, Dr. Redmond O'Connell and Dr. Barbara Hannigan. I wish to express my gratitude to Dr. Richard Roche at Maynooth University for his continued guidance and support since my undergraduate days. I would also like to thank the postdoctoral researchers, research assistants, and fellow postgraduate students in the Whelan Lab over the past many years. There were many coffees and corridor chats that kept me sane on several occasions! Thank you also to the administrative and technical staff at the Trinity College Institute of Neuroscience.

I wish to thank the research participants who contributed their time, energy, and enthusiasm to the research presented here. Finally, I would like to thank my friends and family, and in particular my Mum, Dad, and my brother Christopher for their encouragement and company during what has been an eventful and rewarding journey.

Acronyms & Abbreviations

AIC	Akaike Information Criteria
AUC	Area under the curve
AUDIT	Alcohol Use Disorders Identification Test
BIS	Barratt Impulsiveness Scale
DA	Dopamine
DDM	Drift-diffusion Model
EEG	Electroencephalography
EN	Elastic Net
fMRI	Functional Magnetic Resonance Imaging
FRN	Feedback-related Negativity
ML	Machine Learning
MCMC	Markov Chain Monte Carlo
MMT	Methadone Maintenance Therapy
MRI	Magnetic Resonance Imaging
PE	Prediction error
PFC	Prefrontal cortex
PST	Probabilistic Selection Task
PSTNFB	Probabilistic Selection Task – No Feedback phase
PSTWFB	Probabilistic Selection Task – With Feedback phase
RDoC	Research Domain Criteria
RL	Reinforcement Learning
RLDDM	Reinforcement learning drift-diffusion model
ROC	Receiver Operating Characteristic Curve
RPE	Reward prediction error
STN	Sub-thalamic nucleus

Chapter 1: Literature Review

Substance dependence can be framed as a maladaptive decision-making process, in which substances are persistently sought out by an individual despite negative repercussions. The ability to update reward and punishment contingencies is a fundamental aspect of effective decision-making, requiring the ability to successfully adapt to the changing demands of one's environment. In this programme of research, four experimental studies investigating value-based decision-making in substance misuse and dependence are presented. Study 1 aims to identify individual differences in decision-making between low and high-risk alcohol-misuse groups, Study 2 focuses on a sample of former opioid-users currently in methadone maintenance therapy, Study 3 is a cross-sectional study comparing current and former nicotine dependent groups, and Study 4 extends this by testing if computational, behavioural, and electrophysiological features of value-based decision-making longitudinally predict likelihood of abstinence versus relapse following a smoking quit attempt.

This chapter will provide an introduction to substance dependence and its impact societally, followed by an overview of its neurobiological basis and advances in current theoretical understanding. Next, decision-making research applied to the field of substance dependence will be discussed, with a focus on computational models applied to value-based decision tasks. Machine learning methods and their applications to the field of cognitive neuroscience and substance dependence will also be outlined. Finally, the aims of the current research thesis will be stated.

1. 1. Perspectives on substance dependence

Substance abuse refers to the harmful use of psychoactive substances, including alcohol, nicotine, caffeine, and drugs that affect cognition (e.g., perception, memory, and attention) when consumed. In the European Union, approximately 29% of adults are reported to have tried illicit drugs in their lifetime (European Monitoring Centre for Drugs and Addiction, 2019). Globally, one in every eighteen people have used drugs at least once in the previous year (corresponding to 5.5% of the global population aged 15-65), and 43% of the population have

consumed alcohol in this time (United Nations Office on Drugs and Crime, 2019; World Health Organisation, 2018).

Alcohol use alone is associated with approximately three million deaths each year, with drug-use attributable to 452,000 deaths (World Drug Report, 2018). Tobacco smoking leads to approximately eight million deaths each year, with 1.2 million of these attributable to second-hand or passive smoke exposure (World Health Organisation, 2018). In 2015, more than one in ten deaths were caused by smoking, in addition to 148.6 million disability-adjusted life years (GBD Tobacco Collaborators, 2017). Drug use is associated with chronic and acute health effects, such as drug-related infectious diseases (e.g., HIV, Hepatitis C contracted through injected drug-use), overdose, and dependence. These may be divided into to harm towards others (e.g., social harms such as crime, family adversity, economic costs), or to the drug-user themselves (e.g., physical drug-related and specific mortality and damage, social harms such as loss of relationships and property, in addition to psychological harm such as drug-related impairment of mental functioning and dependence) (Nutt et al., 2010). Alcohol has been reported as the most overall harmful drug in the U.K, with its negative effects on other users surpassing those of opioids (Nutt, King, & Phillips, 2010). Mental and substance use disorders are reported as the leading cause of years lost to disability worldwide, with drug and alcohol use disorders accounting for approx. 21% of disability-adjusted life years – the second highest, following depressive disorders (Whiteford et al., 2013).

While substance-use disorders make a significant contribution to global disease burden (Vos et al., 2017), only a subset of those who use drugs and alcohol become dependent. Thirteen percent of those reported to have used drugs at least once in the previous year suffer from a drug-use disorder, corresponding to a prevalence rate of 0.71 globally (UNODC, 2019). In Ireland, the prevalence rate of alcohol dependence is 3.8% and alcohol use disorders is 8.5%, however the prevalence of heavy episodic drinking in the population is much higher at 37.8% (WHO, 2018). Nicotine addiction is more likely to develop among those who initiate use early in

life; with approximately 50% of those who begin smoking in adolescence continuing to smoke regularly for the next 15-20 years (WHO, 2010). It has been reported that almost 70% of those who use nicotine will eventually transition to dependence in their lifetime, a rate that is much higher compared to other substances of abuse such as alcohol (22.7% cumulative probability estimate), cocaine (20.9%), and cannabis (8.9%; Lopez-Quintero et al., 2011).

Broadly, drug addiction may be defined as *'a chronically relapsing disorder, characterised by compulsion to seek and take the drug, loss of control in limiting intake, and emergence of a negative emotional state (e.g., dysphoria, anxiety, irritability) when access to the drug is prevented.'* (Koob & Volkow, 2010, p. 760). Moderate to severe substance dependence is often considered as 'addiction' (see Rosenthal & Faris (2019) for an etymological discussion on this term), and this is determined by sets of criteria outlined in the ICD-11 and DSM-V.

The ICD-11 (WHO, 2020) and Diagnostic and Statistical Manual of Mental Disorders (DSM-5; American Psychiatric Association, 2013) diagnostic guidelines for Substance Dependence include items relating to physiological withdrawal from and lack of control regarding substance-use, tolerance, and continued use despite recurrent interpersonal, social, physical or psychological problems exacerbated by the effects of the substance.

A definite diagnosis of dependence should usually be made only if three or more of the following have been present together at some time during the previous year:

- (a) A strong desire or sense of compulsion to take the substance;
- (b) Difficulties in controlling substance-taking behaviour in terms of its onset, termination, or levels of use;
- (c) A physiological withdrawal state (see F1x.3 and F1x.4) when substance use has ceased or been reduced, as evidenced by: the characteristic withdrawal syndrome for the substance; or use of the same (or a closely related) substance with the intention of relieving or avoiding withdrawal symptoms;
- (d) Evidence of tolerance, such that increased doses of the psychoactive substances are required in order to achieve effects originally produced by lower doses (clear examples of this are found in alcohol- and opiate-dependent individuals who may take daily doses sufficient to incapacitate or kill nontolerant users);
- (e) Progressive neglect of alternative pleasures or interests because of psychoactive substance use, increased amount of time necessary to obtain or take the substance or to recover from its effects;
- (f) Persisting with substance use despite clear evidence of overtly harmful consequences, such as harm to the liver through excessive drinking, depressive mood states consequent to periods of heavy substance use, or drug-related impairment of cognitive functioning; efforts should be made to determine that the user was actually, or could be expected to be, aware of the nature and extent of the harm.

The severity of substance use disorder ranges from mild substance abuse to severe substance dependence, assessed by the number of criteria that are endorsed. In both of these diagnostic guidelines, many of the criteria may be considered as directly or indirectly related to abnormalities in decision-making (e.g., persistent drug-use in spite of negative outcomes may be considered as a disorder of value and choice, and 'a strong desire or sense of compulsion to take the substance' may refer to failures of cognitive control). The NIH Research Domain Criteria (RDoC) offers a dimensional approach to understanding mental health and illness that emphasises the need to objectively measure changes in neurobehavioral processes at multiple levels of analysis rather than symptom clusters (Insel et al., 2010). A motivation for this approach is that neurobiological processes do not map directly onto diagnostic criteria of psychiatric disorders, and often show overlap. The RDoC is well suited to examining continuums of symptom dimensions, such as craving and cognitive control, and accounts for comorbidity observed in mental disorders.

There is strong comorbidity between alcohol and tobacco dependence, and early initiation of smoking is considered a risk factor for alcohol and substance use disorders in adulthood. In a U.S sample of alcohol dependent individuals, Le Strat, Ramoz, and Gorwood (2010) found that 48% also reported nicotine dependence. In a U.K population survey, Smith et al. (2011) showed that tobacco use and hazardous drinking were associated with poly-drug use involving a wide-range of illicit substances. These studies suggest that the use of multiple substances does not offer a substitutive function whereby increased use in one substance leads to reductions in the other/s, rather it is associated with increased use across multiple substances. Substance dependence also shows high comorbidity with psychiatric conditions (Dani & Harris, 2005); Cahill et al., 2013). Torrens et al., (2011) found that depression was the most common comorbidity with prevalence rates ranging from 12-80% among substance use disorder patients. The co-occurrence of disorders has been shown to affect clinical severity and is associated with poorer treatment response, poorer prognosis, and increased likelihood of attempted suicide compared to those presenting with one disorder (Hasin & Grant, 2004; European Monitoring Centre for Drugs and Drug Addiction, 2019). Adolescents and young adults who use multiple substances in addition to illicit drugs show higher levels of depression and anxiety and increased psychological distress than those with little or no substance use (Jason et al., 2014).

There are many biological and environmental risk-factors leading to addiction, such as genetic (Hartz & Bierut, 2010; Merikangas et al., 1998), gender, personality (high impulsivity), and education (poor educational achievement) factors, this reflects the multi-causal nature of the condition, which is likely to occur at multiple levels (e.g., biological, computational, behaviour) defying simplistic mechanistic explanations (Kendler et al., 2005; 2008). For example, initiation and early substance use may be more strongly associated with social/environmental factors, with later use linked more with genetic factors (Kendler et al., 2008). The addictive properties of the substance itself also contributes to this risk, with nicotine showing the highest

cumulative probability of transition to dependence in the years following its first use, compared to alcohol, cannabis, and cocaine (Lopez-Quintero et al., 2011).

The chronically relapsing nature of substance dependence presents a challenge for successful treatment. Pharmacological therapies for smoking cessation aim to mitigate craving and withdrawal states and act to block the effects of nicotine on relevant pharmacological pathways, such therapies include nicotine replacement therapy (NRT; e.g., nicotine patches, sprays, tablets), bupropion, and varenicline. Varenicline has been shown to outperform NRT and bupropion, more than doubling the likelihood of smoking cessation compared with placebo (Cahill et al., 2013). Other interventions for smoking cessation include group behaviour therapy programmes, financial incentive schemes, and mobile-phone based supports (Notley et al., 2019; Stead et al., 2017; Whittaker et al., 2019). Opioid-agonist treatment (or substitution therapy, typically with methadone or buprenorphine) has shown efficacy in reducing heroin and cocaine-use, reducing co-morbid diseases, and generally increasing the quality of life among patients (Mattick, Breen, Kimber & Davoli, 2009). Methadone has been shown to 'block' the euphoric effects of heroin, with long-acting pharmacological effects that allow for rehabilitation and engagement with normative daily activities. Despite the availability of these treatments, relapse rates among opioid-users can be as high as 91% following treatment (Smyth et al., 2010), highlighting the need to identify relevant mechanisms for treatment and the development of precise interventions (Ginsburg & Phillips, 2018).

1. 2. Neurobiology of Substance Dependence

Koob and Volkow (2010) propose a conceptual framework for addiction consisting of a three-stage cycle that can be studied experimentally in humans and animals. The cycle includes binge/intoxication, withdrawal/negative affect, and preoccupation/anticipation stages, each with differential neurocircuitry and implications for different functional domains. Neuroadaptions at each stage consist of increased incentive salience, decreased brain reward and increased stress, and compromised executive function. These changes are mediated by

three major neurocircuits; the basal ganglia, extended amygdala, and prefrontal cortex, respectively (Uhl, Koob, & Cable, 2019). In the initial stages of the cycle, drug-use behaviour is positively reinforced by positive social rewards and mood states, and this shifts towards a motivation to avoid negative affect and reduce stress in later stages of the cycle. As a result, the value of non-drug rewards becomes blunted in substance dependence

Structural differences have been observed in many studies comparing the brains of those who misuse substances or are diagnosed with an SUD and healthy controls (see Mackey et al., 2019) for mega-analysis of studies on grey matter volume). Reduced grey matter volume in the medial PFC was found in stimulant drug-users (Ersche, Williams, Robbins, & Bullmore, 2013), reduced medial orbitofrontal cortex thickness in smokers (Kühn, Schubert, & Gallinat, 2010), and structural asymmetry in the nucleus accumbens in alcohol and nicotine-dependent participants compared with non-dependent individuals (Cao et al., 2021). Risk of developing stimulant dependence has been associated with reduced functional connectivity in a ventral medial pre-frontal cortex/orbitofrontal cortex, and ventromedial caudate network, indicating possible functional and cognitive decline in areas relevant for goal-directed learning (Ersche et al., 2020). Sweitzer et al., (2016) found increased connectivity in a range of networks prior to a smoking quit attempt among those who successfully abstained for three weeks compared to those who relapsed.

1. 3. Substance dependence and decision-making

As stated previously, substance dependence can be framed as a maladaptive decision-making process, in which substances are persistently sought out by an individual despite their negative repercussions, and in some cases despite an explicitly stated desire to make alternative choices (Redish et al., 2008). A range of behavioural, electrophysiological, and fMRI studies have highlighted aberrations in how individuals with substance dependence make decisions, and these findings seem to corroborate with theoretical accounts of substance dependence such as incentive salience (Robinson & Berridge, 1993), impulsivity (W. K. Bickel & Marsch, 2001), and

dual-process theories (Robbins & Everitt, 1999). In particular, value-based decisions (i.e., those involving the updating of competing choices via valenced feedback) appear to be altered in substance dependence (Redish, Jensen & Johnson, 2008; (Redish, 2004). Following a review of major theories, Redish et al. (2008) proposed a unified framework of addiction involving ten vulnerabilities or 'failure modes' in decision-making processes. These vulnerabilities are proposed to be differentially affected by substance type (e.g., a relationship between the euphoric effects of opioid use and overvaluation of the expected value of a predicted outcome is proposed, whereas the dysphoric effects of initial nicotine use is unlikely to show a relationship with this vulnerability), and their susceptibility interacts with genetic, developmental, and social contexts. In doing so, Redish et al. (2008) sought to account for the myriad individual experiences of addiction through the emergence and co-existence of these decision-making vulnerabilities throughout different stages of the addiction cycle (however see Ahmed, 2008 and Goudie et al., 2008 for further recommendations incorporating social and psychological factors).

Everitt et al. (2001; 2005; 2018) make the distinction between drug-taking and drug-seeking, two distinct and competing psychological processes that reflect stimulus-response and action-outcome learning. The former behaviour is controlled by the reinforcing properties of the drug, and the latter by drug-associated cues which become incentivised across time – to the extent that cues may eventually become more reinforcing than the drug itself. Drug-seeking may occur over longer periods of time, reflecting a preoccupation with obtaining substances that is noted in diagnostic criteria, and involves aspects of planning and deliberation to meet these goals. Drug-taking on the other hand is described by habitual and compulsive behaviour reflected in repeated and seemingly 'out of control' choices to gain substances. These dual systems have been used to describe the transition from goal-directed to habitual behaviour observed in the addiction cycle; from an intentional motivation to increase future value to one that is governed by past actions.

Impaired reward processing distinguishes those with current addictions from controls, pertaining to alcohol use disorder (Wrase et al., 2007), opiate dependence (Gradin et al., 2014; Huhn et al., 2016), pathological gambling (Romanczuk-Seiferth et al., 2015), and nicotine addiction (Bühler et al., 2010; Peechatka et al., 2015). Relatedly, growing evidence suggests that substance dependence involves alterations to value judgements that influence drug-related choices and decisions (Loganathan & Ho, 2021). For example, diminished ventral striatal response to monetary rewards in the monetary incentive delay task was shown as the best predictor of subsequent drug-use in a large sample of adolescents (Büchel et al., 2017). In the following sections, research concerning decision-making impairments in alcohol, opioid, and nicotine dependence are discussed.

1. 3. 1. Alcohol

Altered reward-related brain activity during decision-making has been observed in alcohol dependence, particularly in the stage of outcome anticipation (Galandra, Basso, Capp, & Canessa, 2018). For example, abnormal signal propagation between the ventral striatum and dorsolateral prefrontal cortex (PFC) in response to rewards was observed in alcohol-dependent patients compared with healthy controls, and showed a relationship with levels of craving. Beylergil et al., (2017) found that alcohol dependent patients (ADPs) showed lower punishment sensitivity in a probabilistic reversal task compared to healthy controls, and that more severely affected ADPs showed greater reward sensitivity than patients with less severe dependence. Alterations in decision-making also appear to extend to non-dependent, but harmful, alcohol-users (Lannoy, Billieux, Dormal, & Maurage, 2019). In a study by Rossiter et al., (2012), those who consumed harmful levels of alcohol showed reduced sensitivity to monetary punishment on a monetary incentive Go/No-Go task compared with non-hazardous alcohol users.

1. 3. 2. Opioids

A number of studies have shown that opioid dependence is associated with higher risk-taking and altered reward processing in gambling and decision-making tasks (Brand et al., 2008). In a

non-drug reward task, Gradin et al., (2014) found that patients on methadone maintenance treatment (MMT) showed reductions in insula activation in the anticipation of loss events, and did not encode the successful avoidance of losses as a reward signal in the ventral striatum compared with healthy controls. In a three-month follow-up study of a treatment-seeking MMT sample, relapsers showed increased heroin-cue-induced craving compared with healthy controls and successful abstiners, with changes in craving showing a relationship with nucleus accumbens activation that may be a predictor of future relapse (Q. Li et al., 2015). However, in a review of research investigating brain-related changes in recovery from opioid and methamphetamine dependence, Stewart, May, and Paulus (2019) highlighted limitations regarding the lack of standard outcome measures for recovery, small sample sizes, and lack of follow-up assessments. A small number of studies have shown behavioural decision-making differences between opioid-dependent patients in varying levels of abstinence. Kriegler et al., (2019) found that patients in maintenance therapy performed better on the Iowa Gambling Task, showed less risky decision-making, and reduced craving in response to drug-related cues compared with those who had recently completed detoxification treatment. Passeti et al., (2011) reported that impaired risky decision-making predicted abstinence status at three-month follow-up among treatment-seeking opiate-dependent patients in a community sample, but not in a residential treatment setting. These results suggest that maintenance therapy may improve decision-making, and interact with types of treatment.

1. 3. 3. Nicotine

Nicotine has also been shown to modulate reward-based learning in both human and rodent studies (e.g., Brody et al., 2004). In never-smokers, a single dose of nicotine increased responsiveness to reward cues, lasting for up to one-week following administration (Barr, Pizzagalli, Culhane, Goff, & Evins, 2008). Nicotine-satiated individuals with depression had increased preference for reward stimuli compared with depressed non-smokers, using a reward learning task that tested a preference for a richer reinforcement schedule (Liverant et al., 2014).

Baker et al. (2011; 2013) employed the probabilistic selection task in a substance dependent sample, including those with a range of substance misuse (e.g., alcohol, cannabis, and nicotine use) and reported that these individuals demonstrated attenuated reward positivity compared to non-dependent groups. More specifically, Baker et al. (2011; 2013) showed that their non-dependent group were significantly better at both learning from rewards and from punishers, compared with their substance-dependent sample.

The findings of Baker et al. (2011; 2013) support models of addiction positing that the desensitization of reward circuits over time is associated with addiction (e.g., Rose et al., 2012; Volkow et al., 2016). Whereas acute nicotine administration amplifies reward learning with respect to non-drug cues, chronic nicotine addicts may have a desensitization of the dopamine (DA) system and concomitant blunted reward sensitivity. Indeed, Fehr et al. (2008) have demonstrated that reduced availability of striatal D₂/D₃ dopamine receptors is associated with nicotine dependence, similar to many other types of addiction. However, as Nestor et al. (2018a) note, this is in contrast to the striatal hyperactivity to non-drug rewards observed in some addiction populations. For example, the reward-focused and impulsive behaviour of selecting smaller immediate monetary rewards over larger delayed rewards is well-established in these populations (e.g., Mitchell, 1999).

Garavan, Brennan, Hester and Whelan (2013) proposed that successful abstinence is characterised both by the restoration of brain function once the neurotoxic effects of the drug abuse diminish, and also by the continued process of abstaining from the drug. Briggs, O'Connor, Jollans, O'Halloran, Dymond and Whelan (2015) found that former- and never-smokers, when compared to current smokers, showed greater cognitive flexibility on the contingency shifting Iowa Gambling Task, indicative of an ability to effectively update shifting reward and punishment contingencies in the task. In a reversal task with reward-punishment contingencies, Butler et al. (2017) investigated performance monitoring in current, ex- and non-smokers, showing that the current smoker group had significantly more reversal errors than either of the

other two groups. These studies support the notion that ex-smokers display similar decision-making processes to non-smokers rather than smokers. However, Nestor et al. (2018b) have shown that ex-smokers demonstrate amplified negative valence monitoring compared with smokers and non-smokers, thought to contribute to their successful abstinence.

In recent years there has been an exponential growth in the availability and use of e-cigarette and vaping products (Soneji et al., 2017), which have become popular in part due to a perceived lack of harm compared with traditional methods of nicotine consumption (Kale, Pickeringm & Cooper, 2020) and as a means to reduce cigarette consumption (Barbeau, Burda, & Siegal, 2013). The rise in prevalence of e-cigarette and vaping products warrants investigations into the correlates and predictors of their use, and how this compares within the context of other substance-use (e.g., heavy drinking, cigarette consumption). However, due to its recent emergence, there have been few studies examining cognitive and decision-making correlates of vaping. A small number of studies have investigated trait (i.e., self-reported) impulsivity and e-cigarette use (e.g., Doran & Tully, 2018), with some showing a relationship with lack of perseveration as measured by the UPPS (Chivers et al., 2016; Spindle et al., 2017). In a young adult sample, Lanza et al., 2020) found that increased self-reported lack of premeditation and lower social anxiety predicted e-cigarette use, however Kale et al. (2020) found no significant differences in self-reported impulsivity between e-cigarette users and non-smokers, but showed lower lack of perseverance compared with current smokers. Additionally, Grant et al., (2019) found higher self-reported impulsivity among e-cigarette users compared with non-users on all sub-scales of the BIS.

In contrast to studies on trait impulsivity, there has been a paucity of research exploring task-based measures of impulsivity and decision-making among e-cigarette/vapers. Two studies have studied delayed discounting and e-cigarette use- Stein et al., (2018) showed that e-cigarette users discounted future rewards more than never-smokers, but reported a small effect size with minimal differences observed with former smokers, and Białaszek et al., (2017) found

steeper discounting of delays between e-cigarette users and never-smokers, but did not observe differences with current smokers.

In sum, a large number of studies have investigated changes in decision-making associated with alcohol, opioid, and nicotine dependence. Although impairments in reward processing, tolerance of risk and certainty, and discounting have been observed in the presence of substance dependence, it is not immediately clear how these are affected by treatment and recovery.

1. 4. Electrophysiology

Electroencephalography (EEG) is a method of recording synchronous cortical activity from electrodes typically placed on the scalp. Scalp-recorded electrical activity is generated by post-synaptic potentials from pyramidal neurons located in the cerebral cortex close to the electrode site, and offers high temporal resolution for investigating neural correlates of rapid cognitive processes. Although the content of the EEG signal is not fully understood (i.e., the underlying microcircuit configurations; Cohen, 2017), decades of research has characterised EEG features associated with decision-making.

A substantial literature using reward-related decision-making tasks (e.g., reversal learning, gambling) and others (e.g., oddball, passive viewing tasks) has shown attenuated electrophysiological signals in response to non-drug related rewards, and neural prioritization toward drug-related cues (Stewart & May, 2016). The use of ERPs such as the P300 have been examined as a biomarker of treatment outcomes (Houston & Schliez, 2018), with a growing literature supporting the use of frequency domain EEG data to examine addiction mechanisms (Harper, Malone, & Iacono, 2018a) and elicit clinically relevant behaviour change (see Pfabigan et al., 2011; Wu & Zhou, 2009; Luijges et al., 2018).

The most commonly studied ERPs in the context of substance dependence include error-related and feedback-related negativity occurring post outcome/feedback presentation, the N200 indexing the monitoring of need to inhibit versus activate a response, the P300 as a

measure of the degree of attentional allocation and processing to task stimuli, and the late positive potential as an index of extended attentional capture (Stewart & May, 2016). These feedback related components are also sensitive to magnitude and valence of outcomes, such as monetary gains and losses in decision-making tasks (Pfabigan et al., 2011; Wu & Zhou, 2009).

Differences in the utilisation of feedback reflected in FRN and P3 components have been observed between substance dependent and control samples. In gambling tasks, alcohol dependent samples have shown reductions in FRN and P3 amplitudes, with increased activation in sensory and motor areas. These FRN findings appear to extend to binge drinkers, however no group differences in P3 amplitudes have been observed (Na et al., 2019; Wahlstrom, 2013). In contrast, comparing abstinent heroin users to healthy controls, Zhao et al., (2017) found increased FRN among the heroin-user group, which the authors suggest may be due to higher motivation for high-risk options or unpredicted positive outcomes among this group. In a study comparing cocaine users with longer versus shorter abstinence, and controls, no difference in FRN amplitudes for wins versus losses in a gambling task were observed in the longer abstinent group (Parvaz et al., 2015). Both cocaine-use groups showed an absence of FRN modulation for unpredicted losses compared with controls, suggesting a possible reduction in learning from negative RPE and environmental outcomes (Parvaz et al., 2015). Potts et al. (2014) found differences in medial frontal negativity in a reward expectation task between current and non-smokers, however no group differences in error-related negativity on a flanker task were observed, although ERN was sensitive to reward and punishment. In a study by Muñoz, Anllo-Vento, Fernandez, Montoya, and Vila, (2012), outcome-related amplitudes were sensitive to abstinence and satiation among current smokers, with increased amplitudes for tobacco-related stimuli during abstinence. However, Seow et al., (2019) found no significant association between ERN and alcohol addiction (and other indicators of psychopathology; N=196), suggesting that positive associations reported elsewhere in the literature may be due to small effect sizes.

Resting-state and task-related brain oscillatory activity has also shown associations with substance dependence. A genome-wide association study by Smit et al. (2018) identified novel genetic variants associated with resting-state oscillatory activity and found that hippocampal expression of GABRA2 was associated with both beta power oscillations and alcohol dependence. This confirms previous research demonstrating a relationship between GABRA2, alcohol and illicit drug-use (Mallard et al., 2018), and reward processing (Villafuerte et al., 2012). Resting state beta-band activity was shown to predict subsequent relapse over a six-month period in a substance dependent sample, and greater oscillations in the fast beta range (19.5 – 39.8 Hz) were observed at baseline in relapsers compared with abstainers (Bauer, 2001). Additionally, resting state theta and slow alpha-band activity were significantly reduced in smoking satiation versus deprivation states (Evans et al., 2015; Sutton et al., 2016). Therefore, oscillatory activity may serve as an intermediary step between genes and behaviours relevant to substance dependence.

Frontal midline theta has been proposed as a biophysical mechanism for cognitive control, reflected in spectral activation across the medial pre-frontal cortex, sharing similarities with a range of EEG components (e.g., ERN, FRN, N2) associated with aspects of control such as minimisation of error, novelty, and stimulus conflict (Cavanagh & Frank, 2014). Theta activation in this area may generate a surprise signal that leads to task-relevant adjustments in behaviour such as learning rates that influence prediction errors. Increases in theta activation have been observed in task conditions with high working memory load, and reflect levels of interference between choices in decision-making tasks such as the Stroop, Flanker, and Go/No-Go tasks (Hsieh & Ranganath, 2014; Nigbur, Ivanova, & Sturmer, 2011; Sauseng et al., 2010). Variances in frontal midline theta have been associated with psychiatric dimensions such as anxiety (see Cavanagh & Shackman (2015) for meta-analysis of this effect). Among those with anxiety, a greater influence of negative prediction errors on learning in a probabilistic selection task coupled with increases in theta generated from dorsal midline premotor structures was shown

to result in increased learning from punishment relative to reward (Cavanagh et al., 2011, 2019). Depression symptoms were better explained by changes to delta-band activation arising from orbitofrontal and ventromedial processes in response to reward.

In the context of substance dependence, among a small sample of heroin users (N=15), Li and Xu (2019) found reduced mean theta activation at the Fz electrode between Go and No-go conditions compared with healthy controls ($p=.04$). Accuracy did not differ significantly between groups, however slower reaction times were observed for Go trials in the opioid group. These results present a first hint that opioid use is associated with attenuations in theta activity. In a prospective study of 824 twins, Harper et al. (2018b) examined the association between frontal theta activation during an Erikson flanker task and alcohol-use disorder during the lifetime (303/824 participants had been diagnosed with an AUD by age 29). Reductions in midline frontal theta during response conflict in the task were significantly associated with AUD, and biometric modelling suggested a genetic contribution to this effect. Differences in the onset and duration of AUD from adolescence to adulthood did not show a significant relationship with theta activation. In a Go/No-Go task with unequal probabilities (80/20%; i.e., requiring increased inhibitory control), Holcomb et al. (2019) observed significantly lower theta activation in frontal, central, and parietal regions on No-Go trials among binge drinkers compared to light drinkers in addition to transient increases in beta power among light drinkers associated with preparation to respond in the task. In a gambling task, Kamarajan et al. (2015) showed reductions in monetary outcome event-related theta across gain and loss conditions for a male sample with alcohol dependence compared with controls, differences that also appear between groups at high versus low familial risk for alcohol dependence (Kamarajan et al., 2015). Reductions in frontal midline theta during decision-making tasks with response conflict also occur as a result of acute alcohol consumption (Beaton et al., 2018).

Variations in task-related/functional alpha activity have been observed in a number of clinical domains, including OCD (Perera, Bailey, Herring, & Fitzgerald, 2019) and compulsivity

(Seow, O'Connell, & Gillan, 2020). Pandey et al., (2016) observed group differences in EEG oscillations across all frequency bands between a male sample with alcohol dependence (N=20) and matched controls. During a Go/No-Go task, the alcohol group showed significantly lower 'slow' (8-9.5 Hz) alpha activation in central and parietal regions on Go trials, and significantly lower 'fast' (10-12.5 Hz) activation in parietal, occipital, and temporal regions on No-Go trials. Additionally, the alcohol dependent group showed significantly lower theta activation in frontal, central, and parietal regions, and reduced delta band activity across all regions during the task (Pandey et al., 2016).

1. 5. Probabilistic Selection Task

The probabilistic selection task (PST; Frank et al., 2004) belongs to a class of decision-making paradigms that measure reinforcement learning through 1) a *training phase* in which reinforcement contingencies are learned, and 2) a *test phase* in which participants choose between stimuli. As the PST uses probabilistic, rather than deterministic learning, the reward and punishment contingencies cannot be determined based on the outcome of one trial. Instead, reinforcement history must be integrated over several trials in order to learn the reward/punishment contingencies of the task.

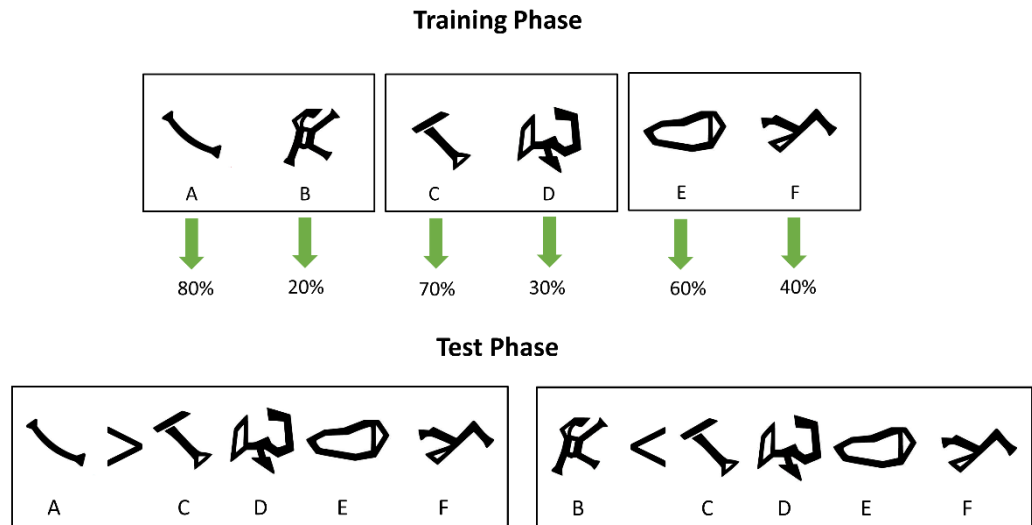


Figure 1. 1. Probabilistic Selection Task (PST) training phase stimulus pairs, with reward contingencies in parentheses, and characterisation of Test phase outcome variables 'Approach A', and 'Avoid B'.

During the training phase, participants learn the reward and punishment probabilities of six stimuli. During the test phase these contingencies are used to guide decision-making. More specifically, during training, three stimulus pairs (denoted AB, CD, EF) are randomly presented with predetermined reward probabilities (A:80%, B:20%, C:70%, D:30%, E:60%, F:40%). Participants are required to pick the most rewarded stimulus in each pair through trial-and-error learning using correct and incorrect feedback. In the Test phase, novel stimulus combinations of the six stimuli are presented without feedback (i.e., a total of 15 novel stimulus pair combinations are possible). Test phase performance is most often quantified by selection frequency of the A stimulus versus the B stimulus in novel pairs. For example, the A stimulus should be preferable to all other stimuli following positive feedback learning, whereas B should always be avoided following negative feedback learning.

The PST may be sensitive to dopaminergic function. In a sample of patients with Parkinson's Disease, Frank et al. (2004) showed that patients on dopaminergic agonist medication (i.e., with sufficient levels of dopamine) learned more effectively from reinforcers than from punishers. The reverse pattern was observed in patients abstaining from medication. Further, in an experiment with healthy older adults (Frank & O'Reilly, 2006), the administration of haloperidol (a DA receptor antagonist) versus cabergoline (a DA receptor agonist) medication affected PST performance, with increased learning from positive feedback observed in the former, and increased learning from negative feedback in the latter. These effects are thought to arise from the selective modulation of striatal D1 and D2 receptors, based on the basal ganglia neural network model of dopaminergic signalling in the striatum proposed by Frank et al. (2004). A number of studies have also shown links between probabilistic reinforcement learning and variations in single-nucleotide polymorphisms in dopamine-related genes, such that expression of DARRP-32 and DRD2 shows a bidirectional relationship with Approach A versus Avoid B selections in the PST test phase (Frank, Moustafa, Haughey, Curran, & Hutchinson, 2007; Frank & Hutchinson, 2009; Sojitra, Lerner, Petok, & Gluck, 2018).

However, it must be noted that a recent study by Grogan et al. (2017) failed to replicate the findings of Frank et al. (2004) in comparing Parkinson's patients on and off medication with healthy controls. Additionally, Maril, Hassin-Baer, Cohen, and Tomer (2013) found that left versus right hemispheric pathology in Parkinson's disease differentially affected reward and punishment learning. This discrepancy in findings may be due to variations in the type of PST used between studies, such as using monetary vs. non-monetary feedback, the discriminability of stimuli used, and the level of accuracy required to pass the training phase (Schutte, Slatger, Collins, Frank, Kennemans, 2017). Overall, these studies highlight that while dopaminergic signalling and the PD hypothesis plays a role in probabilistic reinforcement learning, other factors including working memory, attention, and motivation may also play a large role. For example, Collins, Brown, Gold, Waltz, and Frank (2014) have shown that deficits in working

memory capacity may drive reinforcement learning impairments supposedly due to abnormal striatal dopamine in patients with schizophrenia.

Nonetheless, the PST has been used to examine reinforcement learning in a variety of clinical domains, including schizophrenia (Cicero, Martin, Becker, & Kerns, 2014; Doll et al., 2014; Dowd, Frank, Collins, Gold, & Barch, 2017), depression (Kunisato et al., 2012; Whitmer, Frank, & Gotlib, 2012), attention deficit hyperactive disorder (Frank, Santamaria, O'Reilly, & Willcutt, 2007) obsessive-compulsive disorder (Endrass, Kloft, Klauflmann, & Kathmann, 2011), bipolar disorders (Urošević, Halverson, Youngstrom, & Luciana, 2018), and autism spectrum disorder (Solomon, Frank, & Ragland, 2015). PST performance has shown sensitivity to stimulant medication in ADHD, such that medication improves training phase accuracy, and selectively increases learning from positive feedback in the test phase (Frank et al., 2007). Dowd et al. (2017) reported reduced learning from positive feedback among patients with schizophrenia relative to controls, while Cicero et al. (2014) reported reduced learning from both positive and negative feedback among patients. Conversely, in a comparison of healthy controls and alcohol-dependent patients, Rustemeier et al. (2012) found no group differences in the PST training phase and test phase accuracy, although a near-significant negative relationship between learning from positive feedback and self-reported harm avoidance was observed.

An increasing number of studies have utilised reinforcement modelling in PST research. Cavanagh, Bismark, Frank, and Allen (2019) identified unique dimensions of EEG-related reward and punishment learning during the PST associated with depression and anxiety. Fitting RL models to the training phase, Cavanagh et al. (2019) found that coupling between trial by trial PEs in response to negative feedback and FRN correlated with anxiety, whereas depression was associated with reward-related deficits. Chase et al. (2010) compared learning rates in patients with major depressive disorder versus controls, and showed a 'blunting' effect in patients. That is, reduced learning rates for both positive and negative feedback that correlated with level of

anhedonia. Recently, Baker, Zeighami, Dagher, and Holroyd (2018) showed that reinforcement learning signals were modulated by smoking state in current smokers. More specifically, learning from positive prediction errors increased following cigarette consumption and decreased following abstinence, whereas the reverse pattern was observed for learning from negative prediction errors.

1. 6. Computational Psychiatry

The field of computational psychiatry is based on the assertion that perturbations in how the brain performs computations can lead to psychiatric symptomatology, and aims to translate findings from computational neuroscience into clinical practice (Friston, Stephan, Montague, & Dolan, 2014; Huys, Browning, Paulus, & Frank 2020; Montague, Dolan, Friston, & Dayan, 2012). Research in computational psychiatry has grown rapidly in recent years, with several advances in the prediction of treatment responses (Harlé et al., 2015; Konova et al., 2020), the identification of computational biomarkers relevant to diagnosis (Frässle et al., 2020; Wiecki et al., 2016), and the understanding of mechanisms underlying disorders such as schizophrenia (Collins, Albrecht, Waltz, Gold, & Frank, 2017), autism (Lawson, Mathys, & Rees, 2017), and depression (Cavanagh et al., 2019). For example, Wiecki et al. (2016) used machine learning classification to show that Huntington's Disease state could be predicted in pre-symptomatic individuals (Group A were closer to HD progression as measured by CAG repeat length and age, Group B further were away from disease progression) using a computationally derived parameter of executive function from an anti-saccade conflict task.

In many cases, computational features are more predictive and useful than summary statistics of behavioural task measurements alone. One reason for this is that model parameters reflect an internal process that is made explicit by its function in the model (Huys et al., 2016). Additionally, modelling can capture trial-by-trial variability and decision-making processes that fluctuate within an experimental task by holistically analysing the full range of data— these are often not readily observable by analysing particular aspects of the data (e.g., mean accuracy and

reaction times for conditions of an experimental task). Identifying key latent decision-making parameters that map onto psychological constructs via computational models allows for direct comparisons between tasks and studies- which is not always possible due to methodological variations between tasks. For example, Wiecki (2015) applied a drift-diffusion model to three experimental tasks (numerosity-discrimination task, lexical decision-making task, and memory-recognition task; Ratcliff, Thapar, & McKoon, 2010), and used the estimated parameters from all tasks to create five latent factors predictive of age using a classification technique. The modelling approach also allows for model comparison to formally assess the strengths of competing hypotheses in explaining a given dataset (see Robinson & Chase (2017) for an application of this approach to anhedonia). Finally, the identification of a neurobiological basis for cognitive processes indexed by model parameters may lead to alternative methods of investigation where the use of imaging is not possible due to economic restrictions.

With the steady increase of research into computational psychiatry, efforts have been made to ensure good practices in the use of modelling in cognitive science. This includes developing standards to improve reproducibility, such as making model code and data openly available for other researchers to investigate (Lee, 2018; Lee, Bang, & Kim, 2016; Poldrack et al., 2019), guidelines on how to report modelling results and avoid making incorrect inferences or conclusions from these (Wilson & Collins, 2019) and a framework for the development, validation and deployment of computational assays to improve its ability to address real-world clinical problems (Browning et al., 2020). Lee et al. (2019) suggest (i) pre-registering model predictions and evaluation methods, (ii) post-registering model development and making the models openly available, (iii) conducting detailed evaluation of models to understand their relative strengths and weaknesses (e.g., by bookending models, to include models that are more parsimonious and more complicated than those of primary interest), and (iv) registering model reports. There are multiple levels at which reproducibility must be considered, Benureau and Rougier (2018) suggest that code should be re-runnable (i.e., the execution environment,

software, libraries and dependencies can be recreated), repeatable (e.g., run multiple times by the same researcher with the same results), reproducible (the same results can be obtained independently by another researcher), and reusable (it can be openly used, modified, and improved upon in light of previous research). Taking these steps may limit the effects of flexibility in model analysis on scientific conclusions. Variability in choice of models, methods of estimation, and procedures of inference have been demonstrated to affect conclusions in a blinded multi-team analysis of reaction time data (Dutilh et al. (2019), and elsewhere in cognitive neuroscience (e.g., Botvinik-Nezer et al., (2020) have shown that variability in fMRI analysis workflows resulted in sizable variation of hypothesis results between 70 research teams- no two teams chose the same analysis plan).

1. 6. 1. Reinforcement Learning

The field of reinforcement learning (RL) has been developed from empirical and theoretical advances in many overlapping disciplines (Figure 1. 2), including operant and classical conditioning experiments from psychology (e.g., Hull, 1943; Skinner, 1938), machine learning and reward system research from neuroscience/computer science (Niv, 2009), and optimal control theory from mathematics (Williams, 2009). RL is concerned with how humans, animals and/or artificial agents learn to optimise their behaviour through predicting the perceived consequences of their actions, leading them from one environmental state to the next (Sutton & Barto, 1998).




Figure 1. 2. Diagram of disciplines relevant to reinforcement learning.

Through a process of trial and error learning, state values are updated to increase the probability that more highly rewarded states will be selected in the future to maximise the accumulated reward (Fig. 1. 3). To formalise this, at time-point 't' the agent observes a state 's', takes an action 'a', and receives a reward 'r' as a result of this action. This process is repeated at S_{t+1} and continues along a trajectory of states, actions, and rewards until the termination of agent-environment interaction (e.g., when a game of chess ends). A key feature of RL is the minimisation of error, which is quantified by the prediction error.




Figure 1. 3. Depiction of the agent-environment interface.

The reward prediction error (RPE) signal is proposed to reflect the difference between observed and expected outcomes, and is used to update the value of action-value associations to guide future behaviour (Schultz, 1998); Rescorla & Wagner, 1972). The PE signal reflects expectancies regarding the outcomes of particular choices, and is updated on a trial-by-trial basis in reinforcement learning tasks based on feedback. If a stimulus is consistently paired with a particular outcome, the PE signal is reduced. The PE has gained interest in cognitive neuroscience due to the finding that it is signalled by midbrain dopaminergic neurons (Schultz, Dayan, & Montague, 1997).

There is robust evidence that the functional role of the midbrain dopamine system is to detect and predict future rewards, sending reinforcement learning signals to brain areas involved in decision-making (Schultz, 2002). It has been shown that the updating of stimulus-response values is based on phasic bursts and dips in striatal dopamine which corresponds to reward prediction errors. A seminal study by Schultz et al. (1997) using electroencephalography (EEG) in primates demonstrated that presentation of a conditioned stimulus elicited phasic increases in dopamine cell activity. In contrast, omission of an expected reward led to a dip in dopamine activity during the time window in which the reward was expected (see Figure 1. 4.).




Figure 1. 4. From Schultz et al. (1995). Plots the number of action potentials produced by monitored dopamine neurons within small time intervals during three stages of an experiment. The response of dopamine neurons shifts from initial responses to primary reward (top row) to earlier predictive stimuli (bottom row).

There are a number of algorithms that aim to learn optimal action values in reinforcement learning problems, such as Q-learning, SARSA, and Monte Carlo. These can be dichotomised into model-free and model-based types of learning. Q-learning is an example of model-free learning, which learns action values through trial and error, which has been widely used in the study of trial-by-trial learning in experimental decision-making tasks.

In model-free learning, state-action values are cached based on previous rewards without making representations of future states, an example of this is learning from experience or trial-and-error learning. Actions are selected based on the accumulated average experience of outcomes in the past. This may be considered a form of retrospective learning, whereby future states and actions are expected to behave just as they did in the past. This can lead to

inflexibility in decision-making, as observed in habitual behaviour (Decker, Otto, Daw, & Hartley, 2016). However, this form of learning is computationally inexpensive.

In model-based learning, actions are computed via planning (i.e., an internal representation of future states and actions). According to Dayan and Berridge (2014), a model is an internal representation of stimuli, states, and environmental circumstances that allow prospection. Here, actions are computed based on an internal model of the environment to maximise utility. This allows for flexible updating of reward contingencies in light of future expectations, however comes at the cost of increased computation spent on simulating potential future scenarios. Model-based learning is often related to goal-directed learning (Decker et al., 2016).

1. 6. 2. Computational Psychiatry & Addiction

In this section, a selective review of studies on the computational neuroscience of substance dependence are discussed.

Model-based vs. model-free reinforcement learning

The distinction between model-based versus model-free styles of learning has been proposed to underlie the transition from goal-directed to habitual control observed in addiction. This is typically studied using instrumental ‘two-step’ tasks, or devaluation paradigms translated from the animal literature, such as Pavlovian to instrumental transfer (PIT) or ‘slips of action’ task. The ‘two-step’ task is a sequential decision-making task with probabilistic transitions between two states in each trial that ultimately lead to reward (Daw et al., 2011), and can be modelled computationally to provide a weighting parameter indexing the balance between model-based and model-free learning. Gillan, Kosinski, Whelan, Phelps, and Daw (2016) demonstrated an overarching relationship between reductions in goal-directed learning (indexed by performance in a two-step task) and the symptom dimension of compulsive behaviour and intrusive thoughts, as measured using online questionnaires.

Economic Decisions

Economic choices refer to decisions involving discounting (e.g., delay of smaller immediate rewards relative to larger future rewards) and risks (e.g., loss aversion in gambling tasks and tolerance of uncertainty and ambiguity). Prominent computational models in these types of decisions include hyperbolic discounting, prospect valence learning, and expected utility theory models (Huetzel et al., 2006; Steingroever, Wetzels, & Wagenmakers, 2013; von Neumann and Morgenstern, 1944). The IGT (Bechara et al., 1994) is a decision-making task that requires learning from gains and losses to select the most optimal of four card decks. This requires the management of uncertainty and risk to increase reward. A key finding from the task is that clinical groups tend to persist in selecting disadvantageous card decks throughout the task, whereas non-clinical groups learn to adapt to changes in payoffs by selecting the advantageous decks.

Alcohol

The evidence for a shift from goal-directed to habitual decision-making in alcohol dependence is mixed. Using the two-step task, Nebe et al., (2018) found no association between alcohol use and model-based versus model-free learning in a sample of young male social drinkers. Further, no relationship between alcohol consumption and neural correlates of these relative learning models in the ventral striatum and ventromedial prefrontal cortex was found. In a clinical context, Sebold et al., (2017) compared performance on the two-step task between alcohol-dependent patients and healthy controls. The authors found that alcohol expectancies interacted with model-based control to significantly predict group membership (control, abstainer, and relapse), however model-based control alone did not predict alcohol group status. Relapse was associated with low model-based control and high alcohol expectancies at baseline, rather than low model-based control alone. In an abstinent alcohol dependent sample, Voon et al. (2005) found no significant difference in the two-step task computational model-

free vs. model-based parameter (w) compared with healthy controls. However, significant group differences in learning emerged between binge eating disorder, obsessive compulsive disorder, methamphetamine-dependent samples, and healthy controls.

In a large male sample of alcohol drinkers, Garbusow et al. (2019) found enhanced PIT in high versus low drinkers, and PIT-related amygdala activation which was linked with polygenic risk score for alcohol consumption. In a smaller sample, a similar finding was observed in detoxified alcohol dependent patients, who showed an enhanced PIT effect compared with controls, which was also predictive of alcohol relapse over a three month period (Garbusow et al., 2016).

Opioids

In a series of studies by Ahn and colleagues, computational models were fit to IGT (conceptualised statistically as a four-armed bandit problem) data and fitted parameters were compared between clinical and substance dependent groups. Across three variations of computational models, heroin users showed reduced loss aversion compared with pure amphetamine users and healthy controls, however no overt behavioural IGT differences were observed between the drug-user groups (Ahn et al., 2014).

In a longitudinal study of opiate use disorders in a community treatment setting, Konova et al. (2020) found that a computationally-derived ambiguity tolerance parameter predicted prospective opiate use. Across seven months and 15 experimental sessions, participants completed a lottery task involving choices between guaranteed versus lottery sums of money. The task is designed to measure known-risk versus ambiguity tolerance, and a computational power utility model was fit to parameterise these. At baseline, subjects showed similar performance, however this trajectory changed over time, showing higher ambiguity tolerance associated with relapse. Risk tolerance was not found to predict relapse.

Nicotine

Luijten et al., (2020) found no group differences between smokers and non-smokers in goal-directed versus habitual control, although increased habitual learning was associated with more severely nicotine dependent individuals following appetitive instrumental learning (no association was found following avoidance instrumental learning).

These studies suggest that computational measures may be sensitive to substance-dependence and abstinence. The PST shares similarities with the learning and decision-making tasks highlighted above in terms of the learning mechanisms it seeks to measure. It may be considered a measure of model-free learning, i.e., utilising probabilistic trial-by-trial feedback to increase reward, however its task design (i.e., the lack of a sequential decision-making stage) does not allow for a direct comparison with model-based learning. Importantly, the application of computational models to the PST allows for direct comparisons of latent factors such as learning rates, exploration, and perseverance, which can be applied to a wide range of experimental tasks and may represent task-general markers of substance-dependence.

1. 7. Drift-diffusion Model

Unlike reinforcement learning models, which focus solely on choice behaviour, diffusion models seek to explain the distribution of reaction times that lead to particular choices. The drift diffusion model (DDM) is a form of sequential sampling model that has been extensively used in the analysis of two-alternative forced choice decision-making tasks (Ratcliff & McKoon, 2008). This model posits that decisions involve the gradual accumulation of noisy evidence until a critical decision-threshold is reached and a response is executed. There are four main parameters in the standard DDM (see Figure 1. 5); drift-rate (v), threshold (also referred to as boundary-separation; a), non-decision time (t), and bias (also referred to as starting-point; z). The drift-rate refers to the speed with which evidence accumulation favours one boundary over the other. Boundary-separation provides an index of response caution and the speed/accuracy trade-off. For example, wider boundaries indicate that more evidence is considered before a decision is reached, whereas smaller boundaries indicate faster responding with more noise,

and likely more impulsive decision-making. The non-decision time parameter refers to the time taken to visually encode the stimuli and prepare a motor response. The bias parameter takes into account any pre-existing bias for one stimulus over the other that may influence the decision process. For example, if the starting point is closer one stimulus then less evidence accumulation is required to choose that response. The DDM allows for the comparison of a psychologically meaningful set of parameters between conditions and studies and is sensitive to subtle differences in reaction times between trials. As Moustafa et al. (2015) note, slower reaction times may be due to poor accuracy (measured with drift-rate), slower motor responses (non-decision time), or increased caution (higher boundary separation). These subtleties may be lost in traditional analyses of RT and choice data.

The upper and lower boundaries in Figure 1. 5 refer to each choice on a given trial. In the case of the PST, these could refer to stimulus A versus B in the test phase, with the 'accurate' choice being the stimulus with the highest reward probability from the learning phase (i.e., stimulus A). Wiecki et al. (2013) used a DDM (estimated using a hierarchical Bayesian method) to show that drift-rates in the test phase of the PST varied according to the level of conflict in the stimulus pairs. Lower drift rates were observed for high conflict pairs, which include two stimuli with similar reward probabilities (e.g., A: 80% and C: 70%), compared with low conflict pairs, which included a stimulus with high and low reward probabilities (e.g., A: 80% and D: 30%). In a study using pupillometry, Cavanagh, Wiecki, Kochar and Frank (2014) found that greater pupil dilation predicted increased decision threshold only in high-conflict decisions in the PST.




Figure 1. 5. Diagram of the drift diffusion model, from Wiecki, Sofer, and Frank (2013).

1. 7. 1. Application of DDM

Differences in DDM parameters have been used to distinguish clinical groups from controls. For example, the starting-point parameter has been shown to drive premature responses and impaired response inhibition in patients with schizophrenia compared with controls (Limongi, Bohaterewicz, Nowicka, Plewka, & Friston, 2018). Individuals with high trait anxiety were also shown to increase boundary separation following errors, unlike those with low trait anxiety (White, Ratcliff, Vasey, & McKoon, 2010). Participants diagnosed with autism spectrum disorder were also shown to have wider boundary separation and larger non-decision times compared with controls while performing an orientation discrimination task (Pirrone, Dickinson, Gomez, Stafford, & Milne, 2017).

There is a small but growing literature on the role of DDM decision processes in substance dependence. Two recent studies have shown that acute alcohol intoxication influences particular DDM processes in perceptual decision-making. In a within-subjects design, Stock, Hoffmann, and Beste (2016) showed that alcohol state (sober, intoxicated, or hangover) modulated drift-rate and non-decision rate parameters, while boundary separation was

unaffected. van Ravenzwaaij, Dutilh, and Wagenmakers (2012) showed that alcohol intoxication resulted in increased non-decision times and decreased drift-rates compared with baseline performance in a moving dots task.

1. 7. 2. Combined RL and DDM models

Owing to their development in separate theoretical traditions, DDM and RL models are typically studied independently in the field of value-based decision-making. A number of recent theoretical and empirical studies have highlighted the merits of combining RL with DDMs (Fontanesi et al., 2019; Frank et al., 2015; Militec, Boag, & Forstmann, 2020; Pederson et al., 2017). The main motivation for the integration of these models is to understand how evidence accumulation/drift processes change as learning progresses. For example, evidence accumulation rate towards a response boundary is likely affected by the perceived value associated with a given response. In a probabilistic learning paradigm, the value of responses will change across trials of an experiment. If a given response is consistently rewarded throughout the course of the task, then this will likely result in faster evidence accumulation rates or smaller response boundaries as the task progresses. Therefore RLDDMs may be able to capture the influence of these learning processes on evidence accumulation, response boundaries and starting point biases. Pedersen et al. (2017) have used RLDDMs to analyse the complex effects of medication and clinical status in ADHD, however the application of these models to the field of addiction remains untested.

1. 8. Bridging behavioural and neural data

While computational modelling of behavioural data alone can provide unique insights into the mechanisms underlying reward and punishment learning, further insights may be gained through bridging behavioural and neural data using computational methods. Linking brain and behaviour via cognitive models may improve predictions from otherwise disparate levels of analysis, and a number of methods have been proposed (Turner, Rodriguez, Norcia, McClure & Steyvers, 2016; Turner, Palestro, Miletic, & Forstmann, 2019). Bridwell et al., (2018)

advocate for the use of machine learning to extract single-trial estimates of EEG data that can be linked with behavioural data. For example, Frank et al. (2015) examined the neural mechanisms of reinforcement learning using fMRI and EEG alongside drift-diffusion modelling. The authors found that co-activity between the dorsomedial pre-frontal cortex and subthalamic nucleus was related to fluctuations in the decision-threshold parameter of the DDM according to the level of conflict between response probabilities of stimuli in each pair. Furthermore, Cavanagh et al. (2011) found that increases in trial-by-trial medial pre-frontal cortex EEG theta activity predicted increases in the decision-threshold parameter in high-conflict PST trials. This pattern was reversed with the manipulation of deep brain stimulation of the sub-thalamic nucleus, which resulted in a reduction of decision-threshold and faster reaction times.

1. 9. Machine Learning

As an agnostic, data-driven method, ML is useful in complementing theory- and hypothesis-driven models, and is particularly suited to 'wide' datasets where the number of independent/predictor variables exceeds the number of study participants (e.g., time-series EEG and fMRI data; Obermeyer & Emanuel, 2016). Unsupervised ML models are blind to outcome and dependent variables, therefore the model attempts to cluster feature data into meaningful groups without prior knowledge. Supervised ML models seek to predict a known outcome, such as disease category or relapse status, and often include generalisation tests such as out of sample validation and random label permutation tests to determine their success; if the model can accurately predict the outcome of unseen data, then it is considered successful (Bzdok & Meyer-Lindenberg, 2018). Traditional regression models are subject to 'overfitting', whereby model estimates are optimised for the predictor data that the model is fit to, leading to overoptimistic results that may not generalise to new sample data (Whelan & Garavan, 2014).

The ML approach aims to circumvent such overfitting through the application of resampling/cross-validation and regularisation. Applying a trained model to an external validation dataset that is held separate from the data it was trained on is an ideal method to

ensure generalisability to novel observations. However an efficient cost-effective alternative is the use of cross-validation methods, which involve resampling the data into training and test datasets which are used to generate and apply the model respectively. Resampling procedures include bootstrapping (where random subsets of a dataset are extracted with replacement), leave-one-out cross-validation (where the test dataset contains one datapoint), and k-fold cross-validation (where the dataset is split into k partitions, and the model is generated based on k-1, and applied to the final k partition). For example, the model is first split into a training set (90% of the dataset) and a test set (10%), next the model is fit to the training set, and then the model is evaluated on the remaining test set. In nested cross-validation, a further CV step is implemented within the training set to tune hyper-parameters (such as the regularisation parameters λ and α outlined below). In this case, 81% of the dataset is used to train the model using an array of hyper-parameters and the optimal parameters are evaluated on the remaining 9%. The use of separate datasets to train and evaluate hyper-parameters prevents information 'leaking' into the model and causing overfitting.

Common ML classification algorithms include support vector machines, random forests, and regularised regression. The latter penalises regression weights using regularisation techniques (e.g., LASSO/L1, ridge/L2, Elastic Net) to avoid overfitting caused by overly complex or flexible models. Ridge and LASSO (least absolute shrinkage and selection operator) techniques differ in the type of penalty term they use; ridge adds the squared magnitude of the coefficient as a penalty, and LASSO adds the absolute value of the magnitude of the coefficient. The L1 penalty encourages the exclusion of unimportant predictors in the model by shrinking less important coefficients to 0. This may be useful in identifying significant features in models with numerous predictors. The L2 penalty shrinks predictors to smaller values rather than 0, hence allowing all variables to remain in the final model. The Elastic Net (Zou & Hastie, 2005) is a hybrid technique using a combination of both L1 and L2, with the advantage of allowing

important and correlated predictors to remain in the model. The weight of L1 versus L2 regularisation is controlled by the λ and α parameters.

Therefore, in combining regularisation with cross-validation methods, the generalisability of ML models is improved. One method of assessing ML model performance is to compare it with a baseline classifier using random label permutation. That is, comparing the true model performance with a null model fit to data with randomly assigned outcome labels. This provides an indication of the threshold of significance, and allows for significance testing of ML metrics such as cross-validated r between true and 'null' models.

Machine learning methods are increasingly being harnessed to make individual predictions regarding psychiatric status and clinical outcomes (Dwyer, Falkai, & Koutsouleris, 2018). A systematic review by (Mak, Lee, & Park, 2019) identified 17 studies applying ML methods to addiction research. In the context of alcohol misuse, Whelan et al. (2014) compared the predictive value of various domains (brain (fMRI), personality, cognition, genetics, life history) in predicting binge drinking in an adolescent sample ($n = 692$). 66% of binge drinkers at age 16 were correctly classified using the combined baseline measurements at age 14, personality and life history domains were the most predictive (i.e., highest AUC) individual domains. In a cross-sectional design, Lee et al. (2019) identified 10 features (out of a total 179) that predicted alcohol use disorder treatment-seeking status ($n=778$) with 86% accuracy, and 78% accuracy in an external validation set ($n=236$). In this study, a traditional logistic regression model performed with similar accuracy, however the logistic model required twice as many measurements compared with the ML model. Ahn and Vassileva (2016) also utilised a ML model to classify cocaine-dependent individuals and healthy controls using various task-based measures of impulsivity, with an AUC value ~ 0.90 . These approaches have also been investigated as a tool to guide clinical care and predict patient outcomes (Acion et al., 2017; Connor et al., 2007; Paulus, Tapert, & Schuckit, 2005). Recently, Coughlan, Tegge, Sheffer, and Bickel (2018) used decision trees to identify significant executive function and impulsivity

features to predict smoking cessation treatment outcome. In the field of computational psychiatry, data-driven ML methods may complement theory-driven predictions to investigate latent factors derived from various models and psychopathologies (Gillan et al., 2016; Eisenberg et al., 2019; Wiecki et al., 2015).

1. 10. Unresolved Issues & Current Research

The current chapter has provided an overview of decision-making research relevant to substance dependence, highlighting the role of impaired reward processing in alcohol, opioid, and nicotine dependence. There is a rich literature investigating aberrant decision-making in substance-dependence, however it remains unclear how these are manifested across different substance types and misuse patterns (Ekhtiari, Victor, & Paulus, 2017). As with other branches of psychiatric research, efforts to study the effects of substance misuse on decision-making are sometimes confounded by a strong comorbidity between alcohol, nicotine, and other types of substance-use, which has increased research attention towards trans-diagnostic factors (Eaton, Rodriguez-Seijas, Carragher, & Krueger, 2015). Substance-use is often measured by composite poly-substance outcome variables, and indeed there exist many similarities in the cognitive mechanisms involved across various substance types (e.g., it has been proposed that greater decreases in the valuation of future rewards as measured in delay discounting paradigms is a general marker for substance dependence (Bechara et al., 2019; Bickel et al., 2018), however many others have observed substance-specific effects on decision-making (e.g., Ahn et al., 2014; Peechatka & Janes, 2017). Therefore, the distinct neurobiological and functional aspects of various substances warrant formal comparisons of their effects on decision-making. For example, expectancies regarding opioid-use (immediate euphoria) are likely distinct from cigarette consumption, implicating varying levels of impulsivity and valuation.

Computational models of decision-making provide a quantitative alternative to semantic/verbal outcome variables and summary choice and RT statistics, leading to more direct inferences regarding cognitive components underlying decision-making. An advantage of this

approach is that the influence of various latent variables (e.g., tolerance of uncertainty versus risk, processing speed, learning rate) can be investigated in experimental tasks- this has uncovered subtleties of decision-making and differences between groups that are sometimes not apparent in mean choice and reaction times (Myers et al., 2016). The DDM is a model of decision-making that has been widely applied in the field of cognitive psychology to provide additional insights into decision-making between clinical and normative groups. Concomitantly, RL models have been instrumental in identifying the role of phasic dopamine in signalling RPEs to influence learning and choice behaviour. Recently, DDM and RL models have been combined, allowing for changes in the reward value of task stimuli as they are updated across trials. However, there are a limited number of empirical investigations into value-based decision-making in substance dependence using the DDM of decision-making.

Further, while much research has focused on decision-making differences between groups, less has focused on changes in decision-making across time, and whether successful abstinence remediates impairments observed in former dependence. Some studies have shown that decision-making impairments persist in protracted abstinence, and vary according to substance class (Ahn et al., 2014; Vassileva et al., 2014). While one may expect differences in decision-making between active substance-use, versus short and long-term abstinence, it is not clear how neurobiological mechanisms involved in treatment differentially affect reward processing.

To address these gaps in the literature, I sought to investigate value-based decision-making in substance-dependence and successful abstinence by fitting combined RL DDMs to the PST across four studies. Utilising a machine learning approach with penalised regression models, I compared three sets of features: (i) mean choice accuracy on the PST, (ii) computational parameter estimates, and (iii) self-reported impulsivity as predictors of various substance-use groups.

RLHDDMs were fit to the training phase of the PST, and DDMs to the test phase in each empirical study. Substance-use outcomes were predicted by four models: (i) PST Summary, (ii) Computational, (iii) Personality, and (iv) a Combined feature model, compared with demographic models (i.e., age and gender). In Chapter 3, I sought to predict alcohol misuse risk in a community sample of 115 participants. To investigate the role of abstinence in value-based decision-making, Chapter 4 predicted length of abstinence from heroin among a sample of ex-opioid users (N = 81) receiving methadone maintenance therapy. In Chapter 5, I tested the hypothesis that impairments in value-based decision-making among current smokers would differ from groups of ex-smokers, vapers, and non-smokers (N = 173). Finally, in Chapter 6, changes in value-based decision-making were investigated as predictors of time to nicotine relapse in a longitudinal study tracking individuals (N = 112) through a smoking cessation attempt.

Chapter 2: Methodology

This chapter will provide an overview of the methodology employed across the empirical Chapters 3, 4, 5, and 6.

2. 1. Experimental Task

The PST (Frank et al., 2004) is a two forced choice alternative instrumental learning task designed to test if an individual is more likely to learn from positive relative to negative feedback. The PST comprises a training and test phase. In the training phase, two arbitrary stimuli are presented on each trial and probabilistic positive or negative feedback follows when one of the stimuli are selected. The reward probabilities of the six possible stimuli (i.e., three pairs) are pre-determined (Pair 1: A=80%, B=20%, Pair 2: C=70%,D=30%, Pair 3: E=60%, F=40%). Across 120 trials, the participant learns through trial-and-error to select the stimulus that is most likely correct in a given pair. In the test phase of the PST, each of the six stimuli from the training phase are presented in novel pair combinations without feedback. The participant is required to select what they consider the correct shape in each stimulus pair. The intuition for this phase is that if an individual is more likely to learn from positive feedback, they will consistently select the most rewarded stimuli (i.e., A) in any pair combination. If a person is more likely to learn from negative feedback, they will consistently avoid the most punished stimulus (i.e., B) from the training phase. Here, adopting the method of Cox et al. (2015), positive feedback learning was assessed by calculating the percentage selection of A and C stimuli ($A > CDEF + C > EF$), and negative feedback by the percentage of B and D selections ($B < CDEF + D < EF$). The rationale for this approach is that it is a more sensitive measure than only utilizing Approach A vs. Avoid B outcome variables. The Cox et al. method incorporates more trials, and is likely to avoid ceiling effects where a participant approaches A and avoids B with 100% accuracy.

Two versions of the Probabilistic Selection Task (PST) were employed in this thesis; Version 1 (programmed in Presentation Neurobehavioral Systems, Inc., Berkeley, CA, www.neurobs.com) and Version 2 (programmed by Inquisit, <https://www.millisecond.com>)

were identical in task structure with varying arbitrary stimuli. Version 1 used images from the Stargate science fiction series, whereas Version 2 used Hirigana characters. The use of such arbitrary shapes resulted in a low likelihood that research participants would have a history associated with the task stimuli. Financial incentives for study participation were independent of PST performance across both versions of the task.

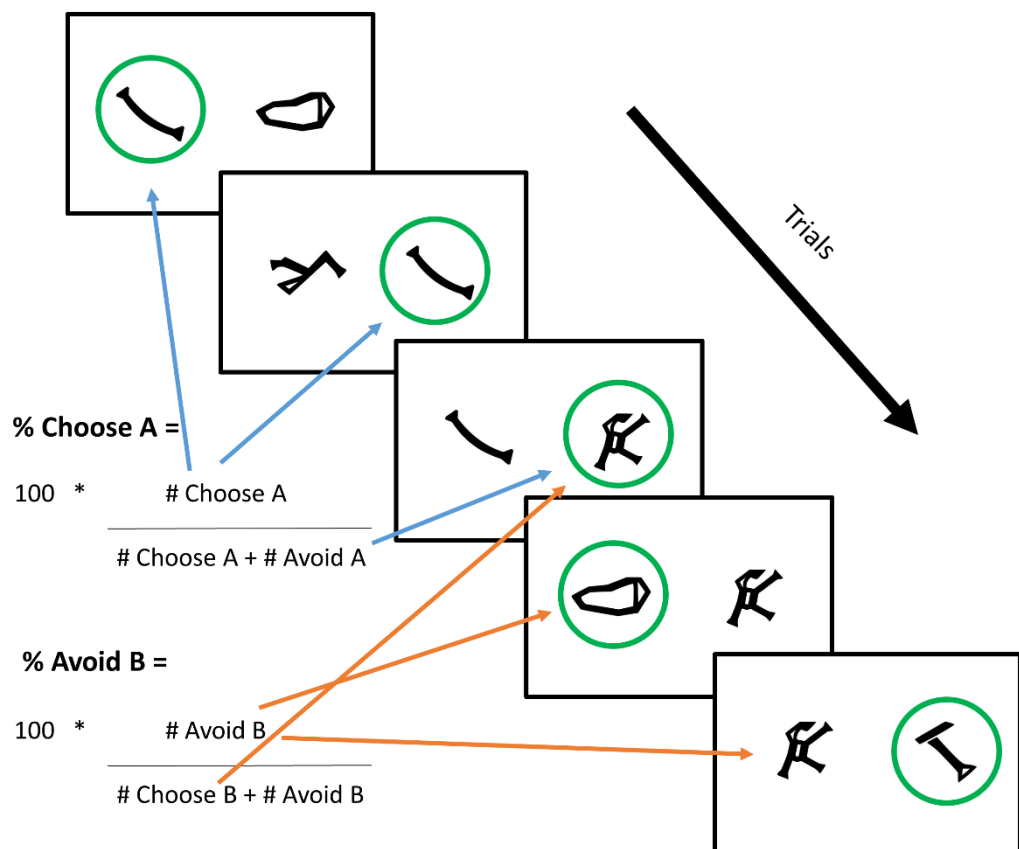


Figure 2. 1. Depiction of how the PST Test phase outcome variables are calculated. The 'Approach A' and 'Avoid B' percentages are calculated from the frequency of A versus. B selections in novel stimulus pair combinations, divided by the sum of total trials in which the A or B stimuli appeared.

2. 1. 1. Probabilistic Selection Task (Presentation)

Presentation (version 20.1) was used to present the PST in Chapters 3, 5, and 6. The stimuli from this version of the task are presented in Chapter 1, Figure 1. 1. In the training phase, participants were instructed as follows:

“You will be presented with two different shapes on the screen.
Pick one. To indicate that you picked the shape on the left,
PRESS THE LEFT BUTTON. To indicate that you picked the shape
on the right, PRESS THE RIGHT BUTTON.

In the practice round you will receive feedback whether your
choice was correct or not.

You may not always receive the same feedback for picking the
same shape, this is part of the task. Just pick whichever shape
you think is correct for each trial.

Please wait for the experimenter to start the task.”

Three pairs of characters from the Stargate series (see Chapter 1, Figure 1.2) were presented randomly on the left and right-hand side of the screen until a response was selected. Feedback in the form of a green ‘✓’ or red ‘X’ was presented for 750 ms, followed by a black fixation cross in the centre of the screen for a duration of 500ms. In the Test phase, participants were verbally instructed to select the shape that they considered correct in each stimulus pair that was presented. Novel pair combinations of the stimuli from the Training phase were presented, the stimuli remained on-screen until a response was selected. A black fixation cross was presented for 500 ms in between trials. A bug in the code resulted in the number of Test phase trials varying between subjects from 96 – 136 trials in Chapter 2, from 90 – 120 trials in Chapter 5, and from 60 to 120 trials in Chapter 6.

2. 1. 2. Probabilistic Selection Task (Inquisit)

The PST in Chapter 4 was presented in Inquisit 4 (<https://www.millisecond.com>). In the training phase three pairs of Hirigana characters (Figure 2.2) were presented randomly on the left and

right-hand side of the screen. Each trial began with an inter-trial interval (250 ms), and ended when a response was selected or 4,000 ms had passed (i.e., non-response trial). Each training phase block consisted of 60 trials, with 20 presentations of each pair. The training phase finished once a participant reached the pre-defined accuracy criteria within a block ($\geq 65\%$ accuracy in AB pairs, $\geq 60\%$ accuracy on CD pairs, and ≥ 50 in EF pairs), or once the participant had completed a maximum of 13 blocks (780 trials). The test phase comprised of novel pair combinations of all six stimuli from the training phase. Participants were required to select the correct stimulus in each pair, however were not provided with feedback.

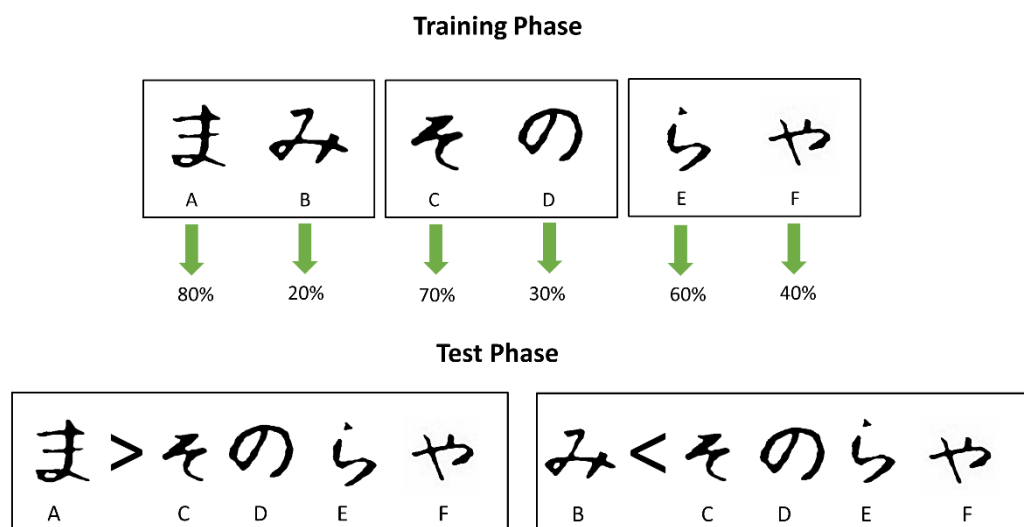


Figure 2.2. Probabilistic Selection Task (PST) presented in Inquisit, training phase stimulus pairs with arrows denoting the probability of reward associated with each stimulus, and characterisation of Test phase outcome variables 'Approach A', and 'Avoid B'.

2. 2 Barratt Impulsivity Questionnaire

Impulsivity is a multidimensional construct that is implicated in substance dependence, and is often measured by self-report questionnaires.

The Barratt Impulsiveness Scale 11th version (BIS-11; Patton & Stanford, 1995) is a 30-item questionnaire designed to measure the personality/behavioural construct of impulsiveness. The

scale is scored with three second-order factors: attentional impulsivity refers to difficulties in maintaining focus on completing a task, motor impulsivity refers to an inability to inhibit responses and tendency to act on the spur of the moment, and non-planning impulsivity refers to present-moment focus with disregard for future consequences. The BIS has shown strong internal consistency and reliability (Stanford et al., 2009), and has been used in clinical and research settings.

2. 3. Fagerstrom Test for Nicotine Dependence

The Fagerstrom Test for Nicotine Dependence (FTND; Heatherton, Kozlowski, Frecker, & Fagerstrom, 1991) was used to measure nicotine dependence in the smoker groups in Chapter 5 and 6. The questionnaire consists of six questions designed to assess physical dependence to nicotine, e.g., ‘How soon after you wake up do you smoke a cigarette?’. Classification of dependence ranges from Very Low, Low, Moderate, High, and Very High.

2. 4. Computational Modelling of PST

2. 4. 1. Hierarchical Drift Diffusion Model

The HDDM package (Versions 0.7.5 and 0.8.0) (Wiecki et al., 2013; Pederson & Frank, 2020) was used to fit drift-diffusion models to trial-by-trial response and reaction time data from both phases of the PST. HDDM utilises hierarchical Bayesian methods with Markov Chain Monte Carlo (MCMC) Slice sampling (Neal, 2003; implemented in PyMC; Patil et al., 2010) to estimate joint posterior parameter distributions, with informed priors from established PST findings in the literature. The HDDM and reinforcement learning HDDM (RLHDDM) used the Wiener first time passage probability distribution (*wfpt*) (Equation 1) to return the probability of choosing *o*, given response time *rt*. Typically, the DDM comprises of starting-point (*z*), non-decision time (*t*), drift-rate (*v*), and boundary separation (*a*) parameters (Ratcliff & Tuerlinckx, 2002). *Starting point* is an index of bias towards one response relative to another. *Non-decision-time* is the time attributed to processing task stimuli and executing a motor response before the decision process itself is executed. *Drift-rate* is an index of the speed and strength of evidence

accumulation until a response boundary is reached. *Boundary separation* reflects the distance between the two response boundaries, wider boundaries require more evidence to be sampled until a response boundary is reached and a decision is executed. Here, a is the boundary separation, t is the non-decision time, z is the starting point, and v is the drift-rate:

$$rt_{o,i} \sim wfpt(a, t, z, v_i) \quad (\text{Eq. 1})$$

In the RLHDDM, the softmax choice rule from traditional reinforcement learning models is replaced with the DDM, with the assumption that the rate of evidence accumulation (i.e., drift-rate) may be described as the scaled difference between the expected value of reinforced options (Pederson & Frank, 2020).

$$v_i = (Q_{upper,i} - Q_{lower,i}) * v \quad (2)$$

In Equation 2, upper and lower Q refer to expected value of the choice options at the upper/lower the bounds of the decision threshold, and v is a free parameter describing exploration/exploitation. Therefore, in the RLHDDM models, drift-rate (v_i) refers to the product of the scaling parameter (v) and the difference between upper and lower Qs on each trial (i).

Individual subject parameters were constrained by a group distribution in HDDMs. The training phase of the PST was fit with a reinforcement learning HDDM (RLHDDM), estimating parameters for drift-rate (v), boundary separation (a), non-decision time (t) and learning rate (with the option of dual positive and negative learning rates, or a singular learning rate). The test phase of the PST was fit with the regular HDDM (estimating drift-rate, boundary separation, and non-decision time). Separate drift-rates were estimated for Win-Win, Win-Loss, and Win-Loss stimulus conditions in the test phase of the PST. The stimulus type referred to the difficulty of determining the correct choice in the novel pair combinations. For example, pairs with two shapes that were more likely to be rewarded in the training phase (e.g., AC, AE, CE) were labelled 'WW', pairs with two stimuli that were more likely to be punished (e.g., BD, BF, DF) were labelled

'LL', and pairs with a combination of highly rewarded and highly punished stimuli (e.g., BC, BE, DB) were labelled 'WL'. The latter stimulus types were considered easier to determine the correct choice compared to 'WW' and 'LL' stimulus types, due to the lower conflict in probability of rewarding/punishing feedback during the training phase. Outliers were handled via HDDM using the 'p_outlier' argument, which found the best-fitting DDM parameters when roughly 5% of data trials could be expected as outliers. Convergence was assessed by running each model three times and comparing variability between chains, generally Gelman-Rubin \hat{r} values <1.2 indicate good convergence. The model trace and autocorrelation plots were also manually inspected to check for convergence of MCMC chains.

Posterior predictive checks were conducted to validate the best-fitting RLHDDM and HDDM models and to identify any systematic inconsistencies between real and simulated data (Gelman et al., 2004). Task data was simulated from the posterior of the fitted RLHDDM model and was compared to the observed data. A random sample from model trace was used to generate accuracy and RT data for each of the training phase stimulus pairs (AB, CD, EF), this step was repeated 50 times to better capture variability in the posterior distribution. These were then compared with the observed data.

Model fit was also assessed using the Deviance Information Criterion (DIC), which measures model complexity by estimating the effective number of parameters, and measures goodness of fit via deviance (Spiegelhalter et al., 2002). Models with lower DICs are considered better supported by the data. The number of samples in each model MCMC chain are reported in the individual chapters, along with DIC comparisons of model fit.

2. 4. 2. Statistical Approach

In each empirical chapter, mean parameter estimates from HDDM and RLHDDM models were used as features in penalised regression models to predict substance-use outcomes. The 'two-step approach' is a key consideration in studies utilising hierarchical computational models of decision-making to explore individual differences. This refers to the process of fitting Bayesian

hierarchical models to groups of subjects separately, and then subjecting individual subject mean parameter estimates to between-groups statistical tests (e.g., ANOVA, t-test). This can cause inflation of effect sizes, as individual-subject estimates shrink towards the group mean in hierarchical models, hence reducing within group variance and increasing the likelihood of a biased statistical finding between groups (Boehm et al., 2018a; 2018b; Moutoussis et al., 2018; Evans & Wagenmakers, 2019). To avoid such biases in the current analysis, HDDM models were fit to all subjects as one group (rather than estimating separate models for high versus low substance-use groups individually). Additionally, identical models were fit to subjects individually to facilitate comparisons with group models, and correlations between parameters estimated from both models are presented in Appendices A, B, and C. Due to convergence issues, the RLHDDMs in Chapter 6 were fit to subjects individually.

A second consideration is the use of so-called ‘point estimates’, such as the mean value of a parameter across multiple MCMC chains, in inferential statistical tests. Hierarchical model point estimates are subject to uncertainty due to shrinkage towards the group mean, and may differ substantially from the true parameter value. Ly et al. (2019) note that this uncertainty must be acknowledged before inferences can be made, and propose the use of ‘plausible values’. These refer to samples drawn from the posterior trace of individual subjects. An effort was made to account for uncertainty associated with mean parameter estimates in the current thesis by conducting ‘robust’ correlations between samples from the group posterior distributions and linear outcome variables (Chapter 3: AUDIT risk-score, Chapter 4: length of abstinence). However, it is acknowledged that further work is needed to fully account for variability in parameter and physiological measures.

2. 5. Machine Learning

Penalised regressions within a machine learning framework were utilised in each empirical chapter. This ML procedure is similar to that described in Kiiski et al. (2018), and an adapted description is presented here (Figure 2. 3).

Predictor variables were first standardised (z-transformed), next the data were split into five cross-validation (CV) folds. The entire analysis was performed five times, using 80% of the dataset (the training set) to create a regression model that was then tested on the remaining 20% of data (the test set). Within the training set, an additional nested CV procedure was conducted to optimise the Elastic Net parameters (i.e., the complexity and weighting parameters, lambda and alpha). Parameters were optimised on 72% of the internal training set, and tested on the remaining internal test set (8%). A range of values was explored to establish the optimal hyperparameters- 15 linearly spaced values of alpha and lambda from 0.01 to 10, and all their possible combinations (i.e., a search grid of 225 parameter-pair values). Results of all five CV folds were aggregated to calculate the frequency with which features were selected from models in different folds.

The entire analysis was repeated 100 times to attenuate the idiosyncrasies of any given model. Results are the mean values across all iterations of the analysis. For logistic models, the outcome metrics were Area Under the Curve and Brier score. For linear models, the outcome metrics were r value and Mean Absolute Error. The performance of each model was further validated by creating null models, which were generated by random label permutation (i.e., shuffled feature data, and fixed covariates). Using the permuted feature data, the entire analysis was performed again. Model accuracy was compared between the true and null models by ranking the cross-validated outcome metrics from iterations of both true and null models, providing an estimate of the level of optimism in the model.

Throughout the empirical Chapters 3, 4, and 5, four models with different sets of features were compared: A PST Summary Model, a Computational Model, a Personality Model, and a Combined Model. Each model included demographic features (age and gender) as covariates, and the null models (i.e., 'Demographic' models) included these as fixed covariates with shuffled non-demographic features. In Chapter 6, a penalised Cox regression was

introduced, using a similar ML procedure, to predict time to nicotine relapse using categorical and continuous time-dependent features. The survival model included censored outcome data (i.e., cases where the true time to nicotine relapse event is unknown, where the event was not observed within the follow-up periods of a study). The Cox models utilised the same Elastic Net procedure, and prediction accuracy was assessed using the Concordance-index (C-index). The C-index, a generalization of the area under the receiver operating characteristic (ROC) curve, is the probability of concordance between observed and predicted survival based on pairs of individuals. Demographic models in Chapter 6 were constructed with age and gender as fixed covariates, and time-dependent features with random data. The accuracy of the demographic models were compared with the test models by comparing mean C-index values across iterations, and were deemed significant where $p < .05$.

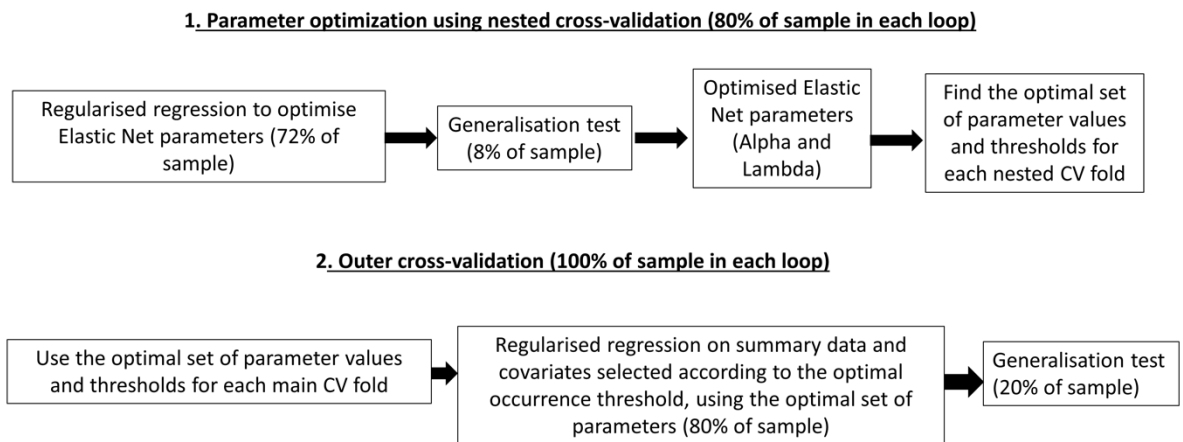


Figure 2.3. Schematic of machine learning method.

Chapter 3: Value-based decision-making as a predictor of hazardous alcohol-use

3. 1. Introduction

Alcohol use is the leading global risk-factor for deaths and disability-adjusted life years among those between the ages of 15-49 (Griswold et al., 2018). Despite the health risks, heavy episodic drinking remains prevalent (16.5% of European population), particularly among 15-19 year olds (31% prevalence rate; Inchley et al., 2014). The deleterious effects of this pattern of drinking remain present even with low-to-moderate levels of drinking on average during the lifetime (World Health Organisation, 2018).

The ability to flexibly update decisions according to changing reward and punishment contingencies in one's environment is one aspect of decision-making that is altered in addiction (Beylegil et al., 2017; de Ruiter et al., 2009), and can be examined by the formal application of computational models. Reinforcement learning models of decision-making often purport that representations of the value of action-outcome associations are updated based on the difference between expected and actual outcomes (i.e., *prediction errors*: PE) following choices (Sutton & Barto, 2015). Midbrain dopamine activity (DA) complies with computational models of reinforcement learning, such that increases in DA are associated with (i) unexpected rewards, (ii) the transition to reward-related cues, and (iii) dips in DA are associated with the removal of reward (Keiflan & Janak, 2015). In other words, the direction and magnitude of dopamine neuron firing is modulated by the expectation of rewards (Schultz, 1998).

Aberrant PEs have been observed in addiction. For example, using probabilistic reversal learning paradigms, substance dependence was characterised by impairments in updating reward values following contingency reversal, with increased perseveration towards previously rewarded responses (de Ruiter et al., 2009; Ersche et al., 2011). Studies utilising reinforcement learning models of decision-making have also shown differential sensitivity to rewards versus punishments in substance dependence. Alcohol dependent patients had lower punishment

sensitivity in a probabilistic reversal task compared to healthy controls, demonstrating a reduced association between PEs and BOLD activity in the dorsolateral prefrontal cortex during (Beylegil et al., 2017).

Decision making also involves several higher order cognitive processes, which can be examined using formal computational models. The drift-diffusion model (DDM) conceptualises decision-making as a process of noisy evidence accumulation via latent psychological processes that are reflected in accuracy and response times (Smith & Ratcliff, 2004; see Chapter 2, Section 2. 4. 1). There is a small but growing application of DDMs to substance-use. Two studies have shown that acute alcohol intoxication influences particular DDM processes in perceptual decision-making (Stock, Hoffmann & Beste, 2016; vanRavenzwaaij et al., 2012). Reduced efficiency of evidence accumulation (derived from drift-rate in a Go/No-Go task) prospectively predicted substance use in a longitudinal study by Weigard et al. (2021), and correlated positively with a component of error-related activation in salience network structures. However, in a study examining the relationship between several experimental task-based factors and real-world outcomes, Eisenberg et al. (2019) found no relationship between binge and problem-drinking, and factors derived from DDM parameters such as response caution. Similarly, Hedge, Powell, Bompas, and Sumner (2020) report no relationship between self-reported impulsivity and response caution.

Computational models of decision-making can detect differences in clinically relevant reinforcement mechanisms that are not apparent in summary data such as mean accuracy and response time (RT) data (Myers et al., 2016; Gueguen, Schweitzer, & Konova, 2021; Huys et al., 2020). This may be particularly important for those individuals at risk of alcohol misuse (cf. substance dependence) with relatively subtle differences. Evidence to date has been mixed. For example, in a sample of young male social drinkers, Nebe et al. (2017) found no association between habitual versus goal directed decision-making and alcohol-use at both the

computational and neural level. However, imbalances in model-based and model-free control at baseline predicted drinking trajectories during a three-year follow-up period (Chen et al., 2021). Cao et al., (2021) also showed aberrations in electrophysiological correlates of reward PEs in high versus low alcohol-use groups, which were not apparent in behavioural task performance alone.

Here, I aimed to test if computational parameters derived from a probabilistic decision-making task (Frank et al., 2004) were associated with alcohol misuse risk. In order to interrogate several variables, I employed a machine learning approach with out-of-sample validation. I compared results to a 'baseline' demographic model, summary PST scores, and with self-report assays of impulsivity. I predicted that the high alcohol group would show reduced learning from punishment compared with reward. It was also anticipated that the high-risk alcohol group would show decreased response caution (indicated by smaller boundary separation in the DDM). However, given the mixed results in value-based decision-making differences among non-clinical individuals to date (e.g., Nebe et al., 2017), specific predictions regarding the direction and significance of particular computational parameters were not made.

3. 2. Methods

3. 2. 1. Participants

Participants were recruited from posters and noticeboards in Trinity College Dublin, and subsequently phone screened for eligibility. All participants provided written informed consent to participate, and the study was approved by the Trinity College Dublin School of Psychology Ethics Committee. Our sample included 82 subjects who were also included in a previous study on impulsivity and alcohol intoxication frequency that included the Probabilistic Selection Task test phase predictor variables (O'Halloran et al., 2018). Participants were provided with €20 in compensation, with up to €10 to reimburse travel costs.

3. 2. 2. Materials

Self-reported alcohol use

The AUDIT is a ten-item questionnaire used to screen alcohol consumption. Items are scored from 0-4 with higher scores indicating higher alcohol use (total possible score of 40). Scores >8 denotes harmful or hazardous drinking, with scores >13 in women and >15 in men, indicating likely alcohol dependence (Liskola, Haravuori, Lindberg et al., 2018; Saunders et al., 1993). The cut-off score of 8 has been recommended in previous studies of alcohol use in college students. Scores on the Alcohol Use Disorders Identification Test (AUDIT) were used to determine low (scores < 8) and high alcohol use (scores \geq 8) groups. 51 low alcohol (Median AUDIT score 5), and 64 high alcohol (Md = 13) participants were identified. Note that the number of alcohol units consumed in the past month was available for 18 out of the 51 subjects in the low alcohol group, and all 64 subjects in the high alcohol consumption group.

Time-line follow-back (TLFB)

The TLFB procedure was used to record a quantitative estimate of alcohol use in the month prior to the testing session. Participants indicated the number and type of alcoholic drinks they consumed on each day of the past month. The TLFB was used to derive (i) the number of monthly drinking days, (ii) the number of binge drinking days in the past month, (iii) the number of consecutive days abstinent, (iv) the highest number of units in one drinking session, and (v) the total number of alcohol units consumed in the past month.

Alcohol Expectancies Questionnaire (AEQ)

The AEQ (Brown, Christiansen, & Goldman, 1987) is a 120 item questionnaire designed to assess anticipated experiences associated with alcohol-use. Six sub-scales assess the domain of alcohol reinforcement: (1) positive global changes in experience, (2) sexual enhancement, (3) social and physical pleasure, (4) assertiveness, (5) relaxation/tension reduction, and (6) arousal/interpersonal power. Each item consists of a statement regarding the effects of alcohol use, and responses are provided on a five-point Likert scale (Disagree Strongly – Agree Strongly).

3. 2. 3. Procedure

Participants completed an online questionnaire via the SurveyCTO (<https://www.surveyccto.com/>) platform within the week prior to the testing session. The testing session took place in Trinity College Dublin, and lasted approximately two hours. Participants completed six computer-based tasks and a Timeline Follow-back (TLFB) questionnaire regarding their alcohol use in the past month. Participants provided written informed consent to participate in the study, which was approved by the Trinity College Dublin Ethics Committee.

3. 2. 4. Data Analysis

Mean choice accuracy on the training and test phase of the PST, and mean reaction times (RTs) between correct and incorrect responses were compared using Bayesian Mann Whitney U tests using JASP software, where values < 0.3 supported the null hypothesis and values > 3 supported the alternative hypothesis (Version 0.11.1; JASP Team 2019). RTs < 0.15 and > 10.0 seconds were removed from the data as outliers (69/13,778 trials in the training phase, 120/12,736 trials in the test phase). Quantile probability plots were used to demonstrate accuracy and RTs across conditions and groups for correct and incorrect responses in the PST (see Ratcliff & Smith, 2011).

Single trial data from the PST were fit with drift-diffusion models using HDDM and parameters estimated from these models were entered into penalised regression machine learning models to predict alcohol-risk, with regularisation using the Elastic Net. Two sets of machine learning models were conducted; (i) logistic regression to predict alcohol risk group and (ii) linear regression to predict AUDIT total score. In each set, four models with separate features were compared; Model 1 (PST Summary) included the mean accuracy variables from the PST, Model 2 (Personality) included the 2nd order BIS scales, Model 3 (Computational) included the computational parameters derived from drift-diffusion models applied to the PST data, and Model 4 (Combined) included both the computational models and the BIS. Age and

gender were entered as covariates in each model. Logistic model fit was assessed using the Brier and AUC scores, and linear models with the MAE and r scores.

A supplementary analysis with residualised features that regress out the confounding effects of covariates is presented in Appendix A, however following the recommendations of Dinga et al. (2020) the current approach was favoured to reduce potential false-positive findings due to information from the residual features leaking into the machine learning model. The regression model features were first standardised using z-scores and then Winsorized whereby any value $>|3|$ standard deviations away from the mean of a given feature was replaced with its maximum value (i.e., $+/-3$). Next, the dataset was divided into five cross-validation folds. Data from each fold was split into training and test sets (80%/20%). Two regression models are reported; (i) a logistic model with high versus low alcohol group as the dependent variable, and (ii) a linear regression model with AUDIT score as the dependent variable. Each model was compared with a baseline 'Demographic Model', which included age and gender as fixed covariates and shuffled data for the remaining features. Model accuracy for (i) was assessed by ranking the cross-validated Brier and AUC scores between demographic and actual model iterations, (ii) was assessed in a similar fashion using the cross-validated r and MSE. The entire analysis was then repeated 100 times with a different random allocation of participants to folds.

Partial correlations were conducted to explore the relationship between model parameters, and the AUDIT, BIS, and AEQ scores, while controlling for age and gender. Parameter values for each subject were sampled from a random chain in the model trace and correlated with each of the questionnaire variables using spearman's r. This step was repeated 1,000 times and the distribution of rho and p-values for each of the parameters and questionnaire variables were plotted across all the iterations. Taking this 'plausible values' approach, I sought to account for variance in the posterior distribution of participant-level parameters (Boehm, Marsman, Matzke, & Wagenmakers, 2018; Ly et al., 2017).

3. 3. Results

From a total of 119 participants with matching questionnaire and PST data, four participants were removed who scored 0 in the PSTWFB. The following analysis is conducted on the remaining 115 participants (Mean age 24.1 ($SD = 8.5$); 60 Female). Scores on the Alcohol Use Disorders Identification Test (AUDIT) were used to determine low (scores < 8) and high alcohol use (scores ≥ 8) groups. 51 low alcohol (Median AUDIT score 5), and 64 high alcohol (Md = 13) participants were identified. Note that the number of alcohol units consumed in the past month was available for 18 out of the 51 subjects in the low alcohol group, and all 64 subjects in the high alcohol consumption group.

3. 3. 1. Behavioral Results

The mean choice accuracy from the PST training and test phases are presented in Table 3. 1. Reaction times are presented in Table 3. 2. For all group comparisons, Bayesian independent samples t-tests (5 chains, 100 samples) were conducted. The low alcohol group had significantly longer RTs on correct trials in the PST test phase, however no remaining group differences were observed. Mean responses on the self-report alcohol-use and impulsivity questionnaires are presented in Table 3. 3. The high alcohol group showed significantly higher risk on all AUDIT subscales, and higher levels of impulsivity on the BIS 1st order motor skills, perseverance, and self-control scales, and the 2nd order attentional scale. Of the total sample, 29 participants reported smoking greater than 40 cigarettes in their lifetime, and 47/98 reported using drugs other than those required for medical reasons within the past 12 months.

Table 3.1. Mean and SD of behavioural choice PST data split by group, with significance test Bayesian independent samples t-tests B10(Mann Whitney U; W) with 1,000 samples and five chains.

<i>Variable</i>	<i>Low Alcohol (n = 51)</i>	<i>High Alcohol (n =64)</i>	<i>BF10</i>	<i>Mann Whitney U (W)</i>
<i>Approach A</i>	75.2 (22.99)	80.39 (21.71)	0.49	1907.5
<i>Avoid B</i>	64.04 (22.06)	65.89 (21.34)	0.21	1716.5
<i>Approach AC</i>	72.67 (19.02)	77.17 (19.23)	0.48	1889.0
<i>Approach BD</i>	60.28 (21.03)	62.38 (21.58)	0.22	1747.5
<i>AB Accuracy</i>	79.07 (15.99)	81.11 (14.74)	0.23	1726.5
<i>CD Accuracy</i>	70.17 (16.12)	77.89 (14.77)	2.93	2094.0
<i>EF Accuracy</i>	52.19 (16.01)	55.16 (19.39)	0.31	1828.0

Table 3.2. Mean and SD of reaction time PST data split by group.

Variable	Low Alcohol (n = 51)	High Alcohol (n = 64)	BF10	Mann Whitney U (W)
<i>PST Training Phase</i>				
All trials	1.12 (0.86)	0.96 (0.79)	0.73	1329.0
AB	1.03 (0.88)	0.86 (0.70)	1.58	1283.0
CD	1.09 (0.86)	0.94 (0.78)	0.92	1332.0
EF	1.12 (0.84)	1.09 (0.86)	0.64	1317.0
Incorrect Trials	1.19 (0.91)	1.01 (0.77)	1.54	1304.0
Correct Trials	1.07 (0.82)	0.94 (0.8)	1.54	1304.0
<i>PST Test Phase</i>				
All trials	1.04 (0.52)	0.83 (0.34)	2.43	1237.0
WW	0.97 (0.51)	0.77 (0.32)	2.17	1238.0
WL	0.94 (0.45)	0.74 (0.30)	3.60	1182.0
LL	1.29 (0.78)	1.06 (0.46)	0.69	1399.0
Correct Response	1.01 (0.54)	0.78 (0.30)	4.88*	1172.0
Incorrect Response	1.23 (0.73)	1.05 (0.48)	0.43	1448.0
Approach AC	8.30 (3.45)	4.02 (0.46)	6766.52*	476.0
Approach BD	8.62 (3.69)	4.13 (0.58)	12721.93*	522.0

Table 3.3. Mean and SD of questionnaire data for the AUDIT, the BIS, and alcohol-related questions split by group. Bayesian independent samples t-tests result (with Mann Whitney U) are included.

Variable	Low Alcohol (n = 51)	High Alcohol (n = 64)	BF10	Mann Whitney U (W)
AUDIT				
Total	4.7 (2.24)	13.41 (4.2)	320417.66	3264.0*
Dependency	1.25 (0.48)	2.16 (1.5)	15797.17	2820.0*
Hazardous Consumption	3.37 (1.77)	6.87 (1.44)	95237.42	3064.0*
Alcohol related Harm	0.86 (1.08)	4.37 (2.89)	21820.5	2935.0*
Alcohol Problems	1.05 (1.13)	6.53 (3.63)	1203000	3092.0*
TLFB				
Monthly Units	-	83.24 (49.42)	-	
Binge Days past Month	-	6.17 (3.32)	-	
Drinking Days past Month	-	9.12 (3.77)	-	
BIS-11				
2nd Attentional	15.84 (3.62)	17.91 (3.99)	4.46	2100.5*
2nd Motor	22.41 (3.21)	24.09 (4.24)	1.50	1993.0
2nd Non- planning	22.82 (4.72)	24.5 (5.14)	0.79	1908.0

3. 3. 2. Computational Models

Drift-diffusion models were fit to single-trial PST data for all participants, both high and low alcohol, in one group (i.e., hierarchically constrained by the whole group prior), and the parameters estimated from these models were entered into a machine learning model to predict alcohol risk. Correlations between parameter estimates and self-reported alcohol use were conducted using samples drawn from the posterior trace. HDDM model fit was assessed using the deviance information criteria (DIC) and is summarised in Table 1. RLHDDM Model 1

was fit with a dual learning rate (positive and negative), whereas model three was fit with a singular learning rate. RLHDDM Model 2 was fit with a dual learning rate and all parameters were estimated separately for high vs. low alcohol group. In the HDDM models, Model 1 was fit with boundary separation (a), drift-rate (v), and non-decision time (t) free to vary, Model 2 estimated separate drift-rate per stimulus type (Win-Win, Win-Lose, Lose-Lose), and Models 3 and 4 were the same as 1 and 2 except each parameter was estimated separately by alcohol risk group (see Appendix A for between-subjects model analyses).

Table 3.4. Summary of HDDM model characteristics, the best-fitting model () was selected based on the deviance information criteria (DIC), with lower scores indicating better fit. Models that estimated separate between-subjects parameters are denoted in parenthesis in the Parameters column. For example, RLHDDM Model 2 estimated separate parameters for high and low risk groups for a (boundary separation), v (drift-rate), and t (non-decision time).*

<i>Model</i>	<i>DIC</i>	<i>Learning Rate</i>	<i>Parameters</i>	<i>Samples (Burn-in)</i>	<i>Max. rhat</i>
<i>PST Training (RLHDDM)</i>					
<i>Model 1*</i>	27129.09	Dual	a, v, t	35,000 (5,000) Thin = 3	1.12
<i>Model 2</i>	27141.59	Dual	(a, v, t) (High-, Low-risk)	15,000 (3,000) Thin = 3	1.03
<i>Model 3</i>	27405.10	Single	a, v, t	15,000 (3,000) Thin = 3	1.06
<i>PST Test (HDDM)</i>					
<i>Model 1</i>	21619.02	--	a, v, t	12,000 (2,000)	1.02
<i>Model 2*</i>	19574.93	--	a, t, v(LL), v(WL), v(WW)	12,000 (2,000)	1.004
<i>Model 3</i>	21618.02	--	(a, v, t) (High-, Low-risk)	12,000 (2,000)	1.04
<i>Model 4</i>	19572.36	--	(a, t) (High-, Low-risk), v(LL), v(WL), v(WW)	12,000 (2,000)	1.004

Two types of RLHDDM model were fit to the training phase of the PST; Model 1, with five parameters (a, v, t, and dual learning rate for positive and negative feedback) showed better model fit (DIC = 27129.09) compared with Model 3, which included a singular learning rate (DIC = 27405.10). Comparing three of the same models, the Gelman-Rubin values for all parameters in Model 1 were ≤ 1.1 (max. value = 1.12), autocorrelation and trace plots were also inspected and showed adequate convergence of chains.

Two types of HDDM model were also fit to the test phase of the PST. Three MCMCs were run for each type of model with 12,000 samples (2,000 burn-in). Model 1, with three parameters (a, t, v) free to vary, showed poorer fit (DIC= 21619.02, max. rhat = 1.02) than Model 2, which estimated separate drift-rate parameters per stimulus conflict (a, t, v(LL), v(WL), v(WW)) again free to vary for the whole group (DIC= 19574.93, rhat=1.005). Mean parameters from the best-fitting HDDM (Model 2) and RLHDDM models (Model 1) were used in the machine learning analysis.

3.3.2.1. Posterior predictive checks

Figure 3. 1. shows the observed versus simulated RLHDDM RT (averaged over 50 iterations) data for each stimulus type (AB = split_by = 0; CD = split_by = 1; EF = split_by = 2). Error responses are negative on the x axis, and correct responses are positive. The simulated data showed similar RTs to the observed data.

Mean choice selections (1=correct, 0=incorrect) were higher in the observed data compared with the simulated data for AB and CD conditions, especially in the first number of trials. The simulated data predicted higher choice accuracy for some EF pair trials.

Posterior predictive checks were conducted on the HDDM model (500 datasets x stimulus condition difficulty), with summary statistics of the actual data falling within the 95% credible interval of the simulated data. Simulated vs. observed RTs are depicted below and indicate good model quality.

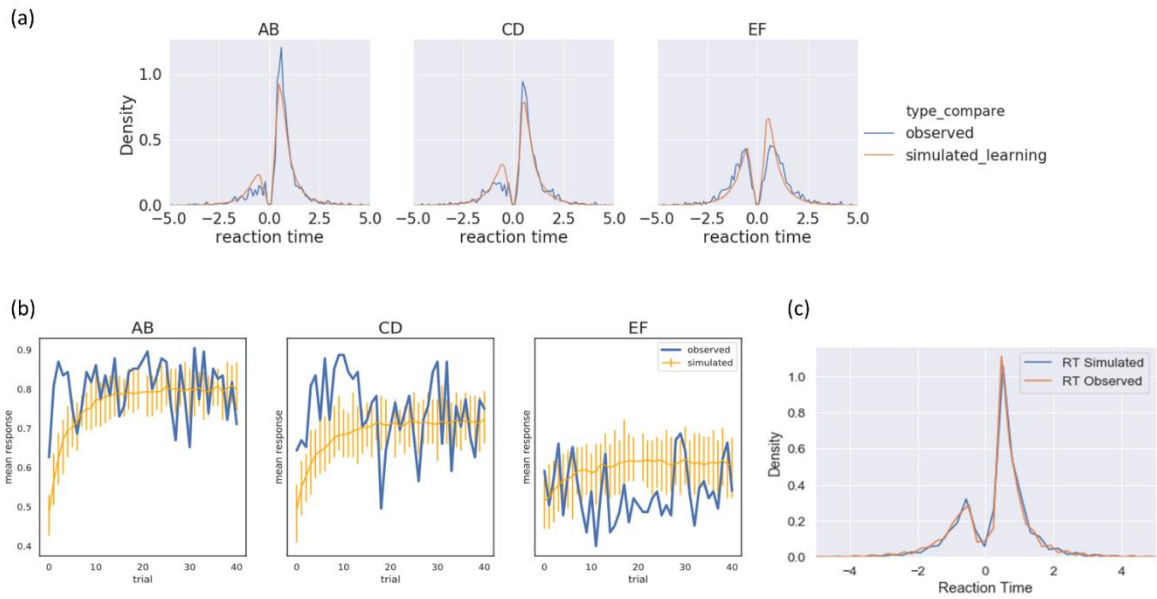


Figure 3.1. Observed vs. simulated reaction times for the RLHDDM (a) and HDDM (c), and response choices in the RLHDDM (b).

Mean model estimates from the best-fitting HDDM and RLHDDM are presented in a series of violin plots below, and mean group comparisons in Table 3. 5. The high-risk group showed significantly higher negative learning rates compared with the low-risk group.

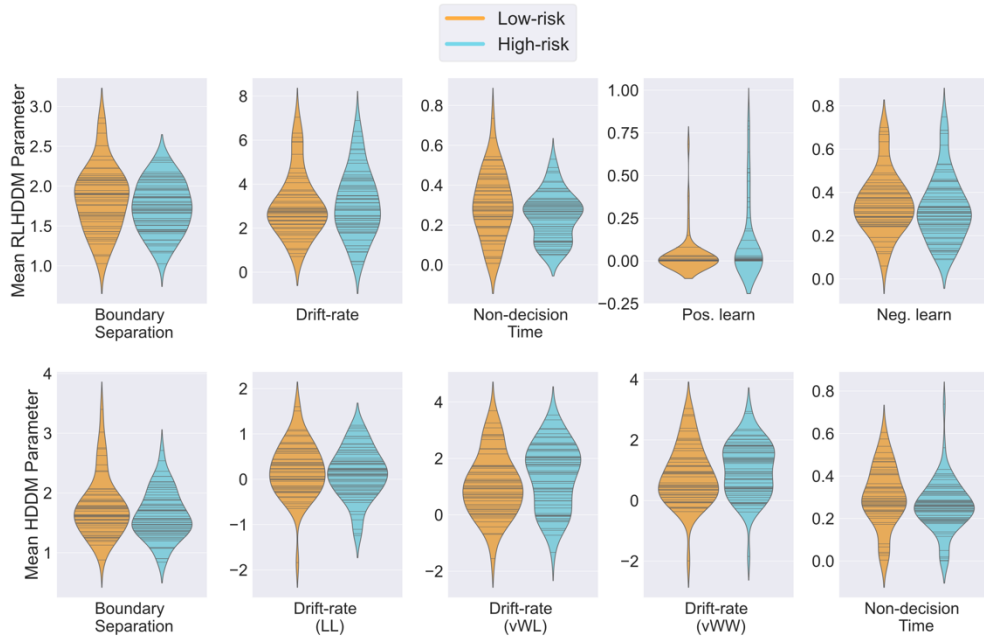


Figure 3.2. Violin plots of (a) HDDM model parameters for the whole group, (b) RLHDDM parameters for the whole group. Risk-group 0 = low-risk, 1=high-risk.

Table 3.5. Mean (SD) of computational parameters with Bayesian Mann Whitney U significance test.

	Low-risk	High-risk	BF₁₀	W
<i>HDDM: Boundary Separation</i>	1.77 (0.50)	1.63 (0.40)	0.504	1410.000
<i>HDDM Drift-rate Lose-Lose</i>	0.21 (0.60)	0.13 (0.56)	0.247	1528.000
<i>HDDM Drift-rate Win-Lose</i>	1.06 (1.12)	1.36 (1.15)	0.368	1891.000
<i>HDDM Drift-rate Win-Win</i>	0.82 (0.96)	1.03 (0.90)	0.315	1860.000
<i>HDDM Non-decision Time</i>	0.31 (0.14)	0.26 (0.12)	1.530	1220.000
<i>RLDDM Boundary Separation</i>	1.84 (0.41)	1.74 (0.33)	0.459	1401.000
<i>RLDDM Drift-rate</i>	3.05 (1.45)	3.12 (1.52)	0.197	1702.000
<i>RLDDM Non-decision Time</i>	0.30 (0.17)	0.24 (0.12)	0.643	1314.000
<i>RLDDM Neg. learning rate</i>	0.03 (0.11)	0.11 (0.22)	14.474*	2235.000
<i>RLDDM Pos. learning rate</i>	0.35 (0.14)	0.33 (0.15)	0.295	1434.000

3. 3. 3. Machine Learning

3. 3. 3. 1. Machine Learning – Alcohol-risk Group Classification

The PST Summary model included the mean approach A/C and B/D from the test phase of the PST as features. The mean model AUC was significantly higher for the original model (0.75; SD = .02) compared with the demographic model (0.73; SD = .02) ($t(198) = 5.82, p < .001$ [.01,.02]) and outperformed in 74% of model iterations. The mean Brier score for the test model (0.194, SD = .01) was not significantly different from the demographic model (0.196, SD = .01), $t(198) = -1.68, p = .09$). Age, gender, and Approach AC survived the 95th percentile of choice frequencies.

Table 3. 6. PST Summary Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Demographic</i>	<i>Mean Beta - Demographic</i>
<i>Age</i>	5	-1.848	5	-1.789
<i>intercept</i>	5	-0.050	5	-0.0258
<i>Gender</i>	4.96	-0.350	4.92	-0.368
<i>Approach AC</i>	4.81	0.271	3.93	-0.029
<i>Approach BD</i>	2.86	-0.021	3.83	0.027

The Computational model included the parameters derived from the computational models of decision-making (boundary separation, drift-rates, non-decision time, and learning rates) as features, with the test model AUC outperforming the demographic model in 85% of iterations. The AUC for the test model (0.76, SD = .02) was significantly higher than the demographic model (0.70, SD = .03); $t(198) = 12.82, p < .001, CI = [0.05, 0.06]$. The test model Brier score (0.20; SD = .01) was also significantly lower than the demographic model (0.21, SD = .01) ($t(198) = -7.96, p < .001$). Age, gender, and all of the computational parameters were significant features.

Table 3. 7. Computational Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Demographic</i>	<i>Mean Beta - Demographic</i>
<i>Age</i>	5	-1.638	5	-1.378
<i>intercept</i>	5	0.060	5	0.081
<i>RLHDDM: Neg. learning rate</i>	4.92	0.367	2.68	-0.002
<i>Gender</i>	4.91	-0.460	4.32	-0.257
<i>HDDM: Boundary Separation</i>	4.7	-0.371	2.63	0.005
<i>HDDM: Drift-rate Win-Lose</i>	4.69	0.583	2.46	-0.027
<i>HDDM: Drift-rate Lose-Lose</i>	4.51	-0.556	2.63	0.012
<i>RLHDDM: Non-decision Time</i>	4.19	-0.176	2.24	-0.006
<i>RLHDDM: Pos. learning rate</i>	4.18	-0.172	2.67	-0.006
<i>RLHDDM: Boundary Separation</i>	3.27	-0.076	2.62	0.01
<i>HDDM: Drift-rate Win-Win</i>	3.09	0.066	2.58	-0.01
<i>HDDM: Non-decision Time</i>	2.92	-0.041	2.52	0.0009
<i>RLHDDM: Drift-rate</i>	2.74	-0.014	2.48	-0.005

The Personality model included the mean 2nd order BIS sub-scales as features. The test model outperformed the demographic models in 71% of iterations based on the mean AUC, and this was significantly higher for the test models (0.75, SD = .01) compared with the demographic model (0.72, SD = .02); $t(198) = 7.64, p < .001$ CI = [.02, .03]. The mean Brier score was significantly lower for the test model (0.19, SD = .01) compared with the demographic model (0.20, SD = .01); $t(198) = -2.88, p = .004$. Age, gender, and the Attentional and Motor BIS scales survived the 95th percentile of choice frequencies, however the Attentional subscale was chosen in almost all model folds on average (4.96).

Table 3. 8. Personality Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Demographic</i>	<i>Mean Beta - Demographic</i>
<i>Age</i>	5	-1.597	5	-1.738
<i>intercept</i>	5	0.043	5	-0.012
<i>BIS: Attentional</i>	4.96	0.320	3.77	0.043
<i>Gender</i>	4.87	-0.372	4.84	-0.352
<i>BIS: Motor</i>	4.25	0.128	3.64	-0.022
<i>BIS: Non-planning</i>	3.75	0.101	3.75	-0.029

In the Combined model, the mean model AUC was significantly higher for the original model (0.74; SD = .027) compared with the demographic model (0.695; SD = .03) ($t(198) = 11.02$, $p < .001$), and outperformed 79% of model iterations. The test model Brier score was significantly lower for the test model (0.20, SD = .01) compared with the demographic model (0.21, SD = .01); $t(198) = -6.39$, $p < .001$. Table 3. 9. below displays the model results with beta weights for each predictor/feature The non-decision time parameter from the HDDM and drift-rate from the RLHDDM did not survive the 95th percentile threshold for significance. Higher negative learning rates in the training phase of the PST therefore is associated with increased odds of being in the high risk alcohol group. This is also the case for attentional impulsivity (i.e., inability to focus attention or concentrate) as measured by the BIS-attentional scale. Lower age was associated with decreased odds of being in the high risk alcohol group.

Table 3. 9. Combined Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Demographic</i>	<i>Mean Beta - Demographic</i>
<i>Age</i>	5	-1.30	5	-1.24
<i>Intercept</i>	5	0.13	5	0.10
<i>BIS: Attentional</i>	4.9	0.27	2.04	-0.0003
<i>Gender</i>	4.88	-0.39	4.13	-0.24
<i>RLHDDM: Neg. learning rate</i>	4.82	0.30	2.2	-0.01
<i>HDDM: Boundary Separation</i>	4.64	-0.35	2.18	-0.001
<i>HDDM: Drift-rate Win-Lose</i>	4.43	0.38	1.74	0.009
<i>RLHDDM: Non-decision Time</i>	4.01	-0.18	1.93	0.002
<i>HDDM: Drift-rate Lose-Lose</i>	3.77	-0.44	1.64	-0.0005
<i>BIS: Motor</i>	3.74	0.09	2.02	0.02
<i>BIS: Non-planning</i>	3.62	0.18	2.18	-0.01
<i>RLHDDM: Boundary Separation</i>	3.03	-0.13	2.08	0.06
<i>RLHDDM: Pos. learning rate</i>	2.9	-0.13	2.29	0.005
<i>HDDM: Drift-rate Win-Win</i>	2.73	0.11	2.02	0.02
<i>HDDM: Non-decision Time</i>	2.65	-0.02	2.07	-0.04
<i>RLHDDM: Drift-rate</i>	1.88	-0.04	2.23	0.02

3. 3. 3. 2. Machine Learning - AUDIT Score

A linear regression was run to predict total AUDIT score. The PST Summary model underperformed the demographic model in 53% of iterations based on r score. The mean score for the test models ($r = 0.41$, $SD = 0.03$) was not significantly higher than the demographic model (0.42 , $SD = .03$); $t(198) = -1.713$, $p = .08$, $CI = [-0.015, 0.0011]$. Similarly, the mean absolute error (MAE) did not differ significantly between test (4.212 , $SD = 0.06$) and demographic models (4.214 , $SD = .06$) ($t(198) = -0.26$, $p = 0.79$). Age was the most significant feature in the model.

Table 3. 10. PST Summary Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Demographic</i>	<i>Mean Beta - Demographic</i>
<i>Age</i>	5	-0.204	5	-0.2047
<i>intercept</i>	5	0.551	5	0.551
<i>Gender</i>	4.67	-0.042	4.64	-0.046
<i>Approach AC</i>	4.23	0.034	2.86	0.0069
<i>Approach BD</i>	1.84	-0.003	3.03	0.0024

The Computational parameter model outperformed the demographic models in 53.9% of iterations. The mean original model MSE(0.363, SD = .037) was significantly lower than the demographic model (0.359, SD = 0.41), $t(1998) = 2.22, p=.03$; $t(1998)= 2.5, p=.01, CI = [-.001 -- .01]$. The mean original model r (0.363, SD = 0.037) was significantly higher than the demographic model (0.359, SD = 0.041); $t(198) = 2.21, p = 0.026$. Mean absolute error for the original model was 4.28 (0.07), and 4.25 (0.09) for the demographic model, this difference was marginally significant ($t(198) = 2.105, p - .036$). The selection of features are presented below, age, intercept, negative learning rate, gender, drift-rate from the RLHDDM, and drift-rate for Win-Losses in the HDDM survived the 95th percentile.

Table 3. 11. Computational Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Demographic</i>	<i>Mean Beta - Demographic</i>
<i>Age</i>	5	-1.913	5	-2.01
<i>Intercept</i>	5	9.412	5	9.39
<i>RLHDDM: Neg. learning rate</i>	4.728	0.714	1.79	0.007
<i>Gender</i>	4.295	-0.449	4.18	-0.42
<i>HDDM: Drift-rate Win-Lose</i>	4.196	0.323	1.37	-0.002
<i>RLHDDM: Drift-rate</i>	2.508	0.270	1.53	0.002
<i>RLHDDM: Boundary Separation</i>	1.75	-0.245	1.46	-0.002
<i>HDDM: Non-decision Time</i>	0.985	-0.153	1.59	0.004
<i>HDDM: Boundary Separation</i>	0.722	-0.180	1.57	-0.008
<i>RLHDDM: Non-decision Time</i>	0.611	-0.123	1.45	0.003
<i>RLHDDM: Pos. learning rate</i>	0.527	-0.143	1.71	0.001
<i>HDDM: Drift-rate Win-Win</i>	0.473	0.115	1.52	-0.007
<i>HDDM: Drift-rate Lose-Lose</i>	0.417	-0.200	1.63	0.003

In the Personality Model, the test model r (0.49, SD = 0.03) outperformed the demographic model r (0.39, SD = 0.03) in 95% of iterations, and was significantly higher ($t(198) = 24.70$, $p < .0001$, CI = [2.69, 3.15]). The MAE for the test model (4.06, SD = .08) was significantly lower than the demographic model (4.21, SD = .07); $t(198) = -14.72$, $p < .001$. All features were selected in > 4 model folds on average, however age was the only feature to survive the 95th percentile threshold.

Table 3. 12. Personality Model mean choice frequencies and beta values for each feature in the test and demographic model.

Feature	Mean choice frequency - Test	Mean Beta - Test	Mean choice frequency - Demographic	Mean Beta - Demographic
<i>Age</i>	5	-2.050	5	-2.171
<i>BIS: Non-planning</i>	5	0.962	2.2	-0.032
<i>intercept</i>	5	9.392	5	9.389
<i>BIS: Motor</i>	4.97	0.736	2.11	0.0269
<i>Gender</i>	4.9	-0.591	4.66	-0.487
<i>BIS: Attentional</i>	4.31	0.327	2.38	-0.0196

In the Combined Model, the original models performed better than the demographic models in 89.5% of iterations (based on mean square error and Pearson’s r). The mean MSE for the original model (-25.19) was significantly higher than the demographic model (-27.68) ($t(1998) = 55.4$, $p < .001$, $CI[-2.4 - 2.6]$). The r value for the original model (0.45, $SD=.034$) was also significantly higher than the demographic model (0.35, $SD=.042$) ($t(1998)=55.79$, $p < .001$, $CI[-.09 - 0.1]$). The MAE for the original model (4.12, $SD = .07$) was significantly lower than the demographic model (4.26, $SD = 0.08$) ; $t(198) = -12.36$, $p < .001$. Beta values for the individual features are presented below: age, intercept, BIS sub-scales, negative learning rate, boundary separation, and drift-rate from the RLHDDM, and drift-rate for win-loss pairs in the HDDM survived the 95th percentile correction.

Table 3. 13. Combined Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Demographic</i>	<i>Mean Beta - Demographic</i>
<i>Age</i>	5	-1.79	5	-1.97
<i>Intercept</i>	5	9.41	5	9.39
<i>BIS Non-planning</i>	4.97	0.83	1.39	-0.005
<i>BIS Motor</i>	4.90	0.67	1.39	-0.009
<i>RLHDDM: Neg. Learning Rate</i>	4.64	0.61	1.66	0.008
<i>Gender</i>	4.57	-0.49	4.04	-0.40
<i>BIS Attentional</i>	3.8	0.32	1.49	0.002
<i>HDDM: Drift-rate Win-Lose</i>	3.675	0.243	1.19	-0.002
<i>RLHDDM: Boundary Separation</i>	2.883	-0.278	1.36	0.002
<i>RLHDDM: Drift-rate</i>	2.668	0.266	1.48	0.008
<i>HDDM: Drift-rate Win-Win</i>	1.551	0.167	1.31	-0.019
<i>HDDM: Boundary Separation</i>	1.115	-0.212	1.52	-0.014
<i>RLHDDM: Non-decision Time</i>	1.049	-0.133	1.39	-0.0004
<i>RLHDDM: Pos. learning</i>	0.772	0.121	1.52	0.007
<i>HDDM: Non-decision Time</i>	0.558	-0.099	1.34	-0.005
<i>HDDM: Drift-rate Lose-Lose</i>	0.357	0.029	1.38	0.0002

Machine learning scores are presented in Figure 3. 3. Overall, age was the most predictive of alcohol use risk and AUDIT score, surviving the 95th percentile threshold in all of the ML models. Of the logistic models predicting alcohol-risk group, the best-fitting model according to AUC was the computational features model (0.76), with the BIS and PST summary models also showing similar fit (0.75). Of the linear models, the 2nd order BIS features provided the best fit ($r = 0.49$). Of the computational decision-making features, negative learning rate from the RLHDDM and

drift-rate for low conflict stimulus pairs in the HDDM were the most frequently selected features across model folds and iterations. Mann Whitney U tests comparing Brier and MAE metrics between the PST Summary, Personality, Computational, and Combined Models showed a significant difference in Brier score between the Combined model and the remaining models (all p 's $<.0001$). Not significant differences were found between mean Brier scores for the remaining models (all p 's $>.24$). Mean MAE for the Personality Model was significantly lower than the remaining models, and the remaining models differed significantly from each other (all p 's $<.0001$).

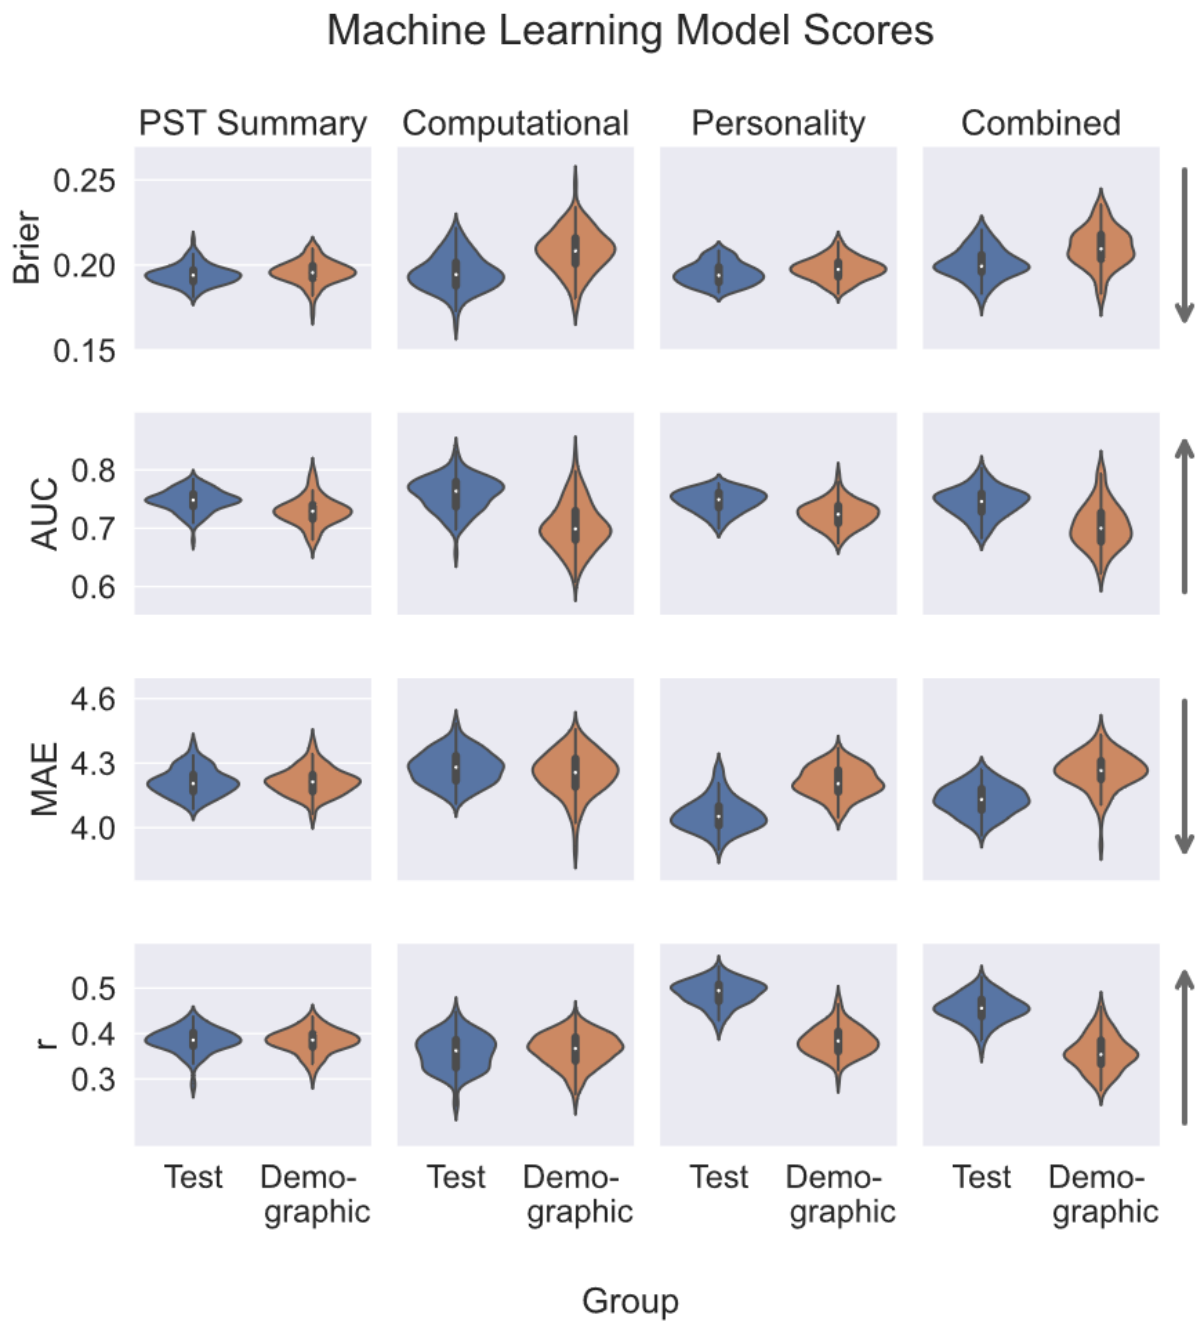



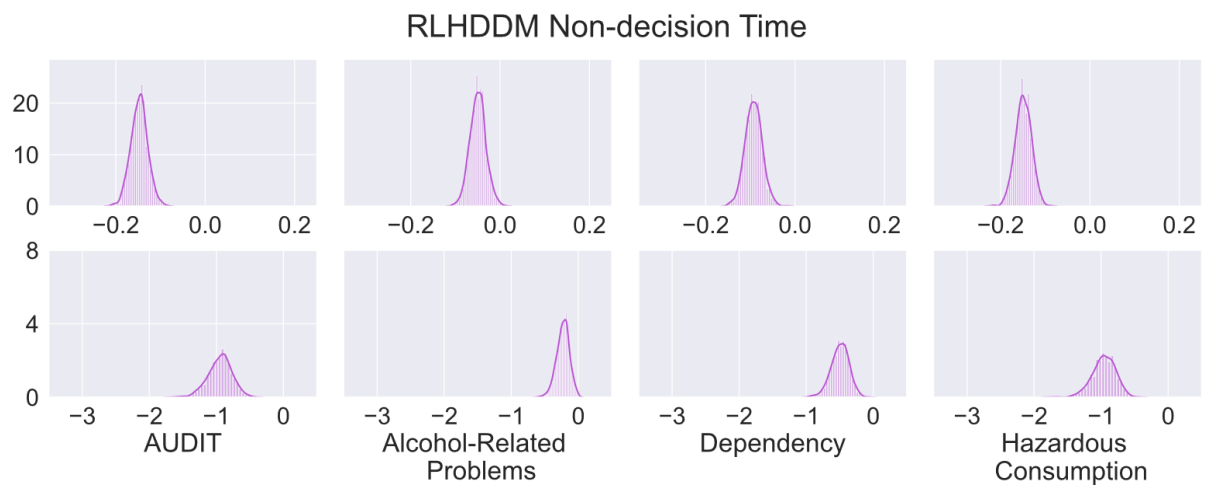


Figure 3.3. Violin plots of mean Brier and AUC for logistic models, and MAE and mean r for linear models. The PST Summary includes features from the PST test phase, the Personality Model included the BIS 2nd order sub-scales, the Computational Model included the (RL)HDDM model parameters, and the Combined Model included the BIS and (RL)HDDM model parameters as features. All models are compared with a Demographic Model, which included age and gender as fixed covariates.



3. 3. 4. Correlations

Robust correlations between AUDIT scores and DDM parameters (Figure 3. 4.) suggest that negative learning rate correlates positively with AUDIT score (i.e., higher alcohol-risk is associated with higher rate of learning from negative outcomes), with the reverse relationship observed for the positive learning rate (i.e., higher alcohol-risk is associated with lower rate of learning from positive outcomes). However, the distribution of significance thresholds suggest that these correlations were not robust. A more robust negative correlation was observed between the AUDIT and the non-decision time parameters from the PST training and test phase, particularly between the total AUDIT score and the Hazardous Consumption scale. Similarly, a negative correlation was found between the boundary separation parameters from the PST training and test phases. The drift-rate parameter for the Win-Loss condition in the PST test phase showed a positive correlation with the total AUDIT score, alcohol-related problems, and hazardous consumption scales.

3. 3. 4. 1. AUDIT Plots







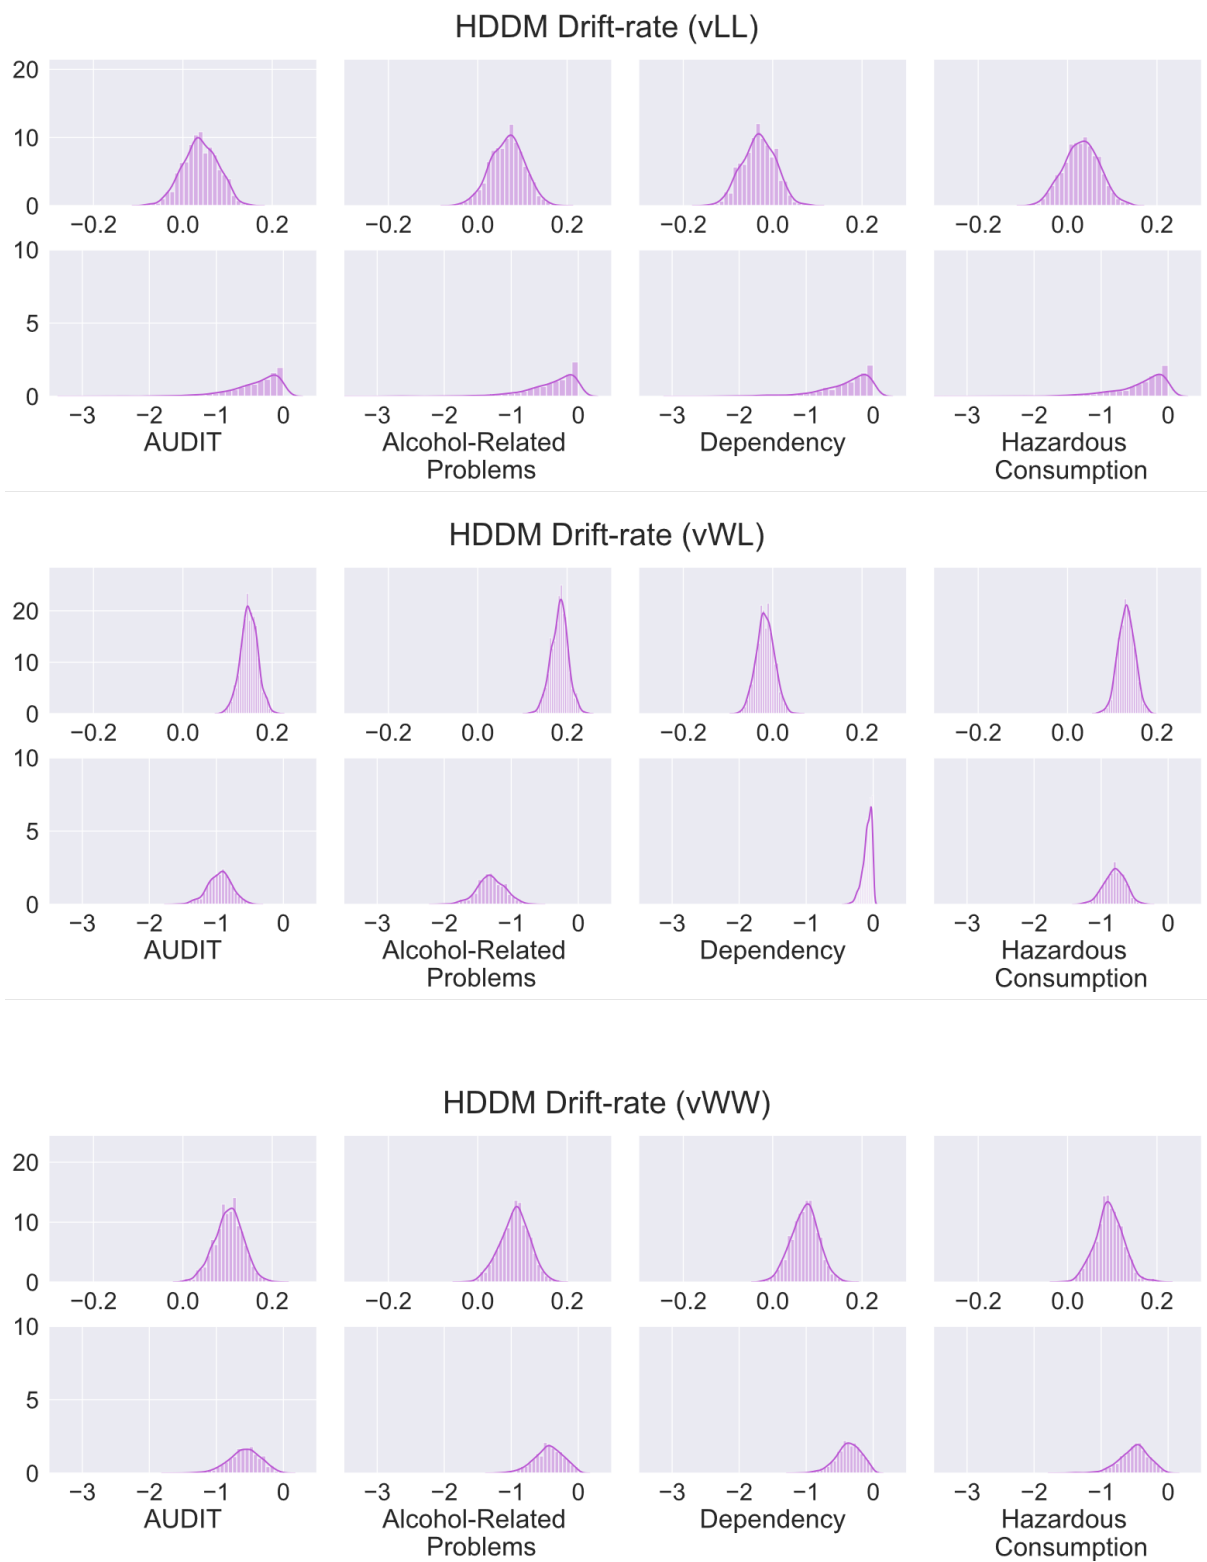


Figure 3.4. Distribution of rho (top row) p-values (bottom row) for 1,000 correlations between HDDM and RLHDDM parameters and questionnaire variables.

Each correlation was compared with 1,000 null models of shuffled AUDIT scores. The distribution of null model rhos centred on 0 and showed greater variance compared to the test correlation. Distributions of rho values for the test and null models were compared by examining the proportion of values in the test models that were less than the null model values. Correlations with more significant rho values had a greater percentage of difference compared with the null models (which centred on 0). These are illustrated for each parameter and AUDIT sub-scale below (Figure 3. 5; null models in orange, test models in blue). For example, the 't' parameter in the RLHDDM had a larger proportion of rho values less than its null model for the AUDIT total score (95.2%), indicating a more robust correlation than that observed for the positive learning rate parameter (65.3%). Binomial cumulative distribution functions were computed for each correlation difference with the probability of success = 0.5 (all probabilities = 1 or 0, indicating that the difference was unlikely to occur by chance).

Test vs. Null Correlations




Figure 3.5. True versus null robust correlation distributions for each RLHDDM and DDM parameter, with percentage of values greater than the null.

3.3.4.2. Correlations between ML features:

Correlations between mean parameter estimates, PST summary scores, and the BIS are presented in Table 3. 14. The BIS sub-scales showed were highly positively correlated with each other, as were many of the computational parameters. Boundary separation and non-decision time from the PST training phase model correlated significantly with the test phase, and drift-rates from the test phase showed a significant relationship with the mean Approach AC/BD summary scores. Positive learning rate correlated negatively with the attention and non-planning sub-scales of the BIS.

Feature	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1. BIS Attention	—														
2. BIS Motor	0.519***	—													
3. BIS Non-planning	0.586***	0.529***	—												
4. RLHDDM: a	0.089	0.023	0.067	—											
5. RLHDDM: v	-0.011	-0.005	0.041	-0.246**	—										
6. RLHDDM: t	0.052	-0.038	0.091	0.319**	0.068	—									
7. RLHDDM: Neg. Learning	-0.001	0.063	0.116	-0.001	-0.109	0.089	—								
8. RLHDDM Pos. Learning	-0.193*	-0.122	-0.189*	-0.069	0.031	-0.019	-0.110	—							
9. HDDM a	0.089	0.102	0.001	0.538***	-0.118	0.290**	0.036	-0.016	—						
10. HDDM vWW	0.026	-0.016	-0.053	-0.134	0.440***	-0.019	0.105	0.266**	-0.105	—					
11. HDDM vLL	0.010	-0.209	-0.062	-0.277**	0.324**	-0.179	0.067	0.070	-0.111	0.467***	—				
12. HDDM vWL	0.099	0.087	-0.044	-0.055	0.530***	0.045	0.189*	0.094	0.089	0.598***	0.553***	—			
13. HDDM: t	0.015	-0.064	-0.014	0.325**	-0.103	0.709***	0.053	-0.018	0.273**	-0.129	-0.237*	-0.028	—		
14. Approach AC	0.095	0.026	0.022	0.048	0.410***	0.136	0.186	0.180	0.144	0.880***	0.395***	0.746***	0.012	—	
15. Approach BD	0.095	0.070	0.012	-0.108	0.387***	-0.045	0.165	0.004	0.111	0.434***	0.805***	0.830***	-0.109	0.534***	—

*p < .05, ** p < .01, *** p < .0001

Table 3. 14. Partial Spearman's correlations between machine learning features, with age as a covariate. Significant correlations are highlighted between (i) self-report measures (orange), (ii) task-related parameters and variables (blue), (iii) task-related and self-report measures (green).

3. 4. Discussion

This study sought to investigate if computationally derived parameters of value-based probabilistic decision-making predict alcohol misuse risk. I fit combined reinforcement learning and drift-diffusion models to trial-by-trial PST data and used parameters derived from these to predict alcohol risk using penalised logistic and linear regression. Additionally, I compared the Computational model to a Personality model including impulsivity features from the BIS, and a PST Summary model including mean accuracy in the PST test phase as features. Learning rate for negative prediction errors was the most frequently selected feature across the Computational machine learning models, followed by drift-rate for the Win-Loss condition in the total AUDIT score models, and boundary separation from the RLHDDM in the AUDIT risk-group classification.

Higher negative learning rates, higher impulsivity, higher drift-rates, lower boundary separation, lower non-decision time, and lower age were all associated with higher alcohol misuse risk. Females had significantly reduced odds of being in the high alcohol risk group. In the Combined feature models for AUDIT risk-group classification, all computational parameters were significant features except the non-decision time parameter from the HDDM, and the drift-rate parameter from the RLHDDM. In the Combined feature model for total AUDIT score, negative learning rate, boundary separation, and drift-rate from the RLHDDM, and drift-rate (WL) from the HDDM were significant features. The impulsivity features were also significant in these models. These findings were supported by correlations between the AUDIT and mean computational parameters sampled from the posterior, which showed robust negative correlations between boundary separation, non-decision-time, and total AUDIT score. Positive correlations were observed between the drift-rate for Win-Losses and the AUDIT.

We found that increased sensitivity to negative prediction errors (i.e., higher negative learning rates) is associated with increased alcohol misuse. However a number of studies have shown *reduced* sensitivity to losses in substance dependence (e.g., Beyergil et al., 2017; Tanabe et al., 2013; Vanderschuren, Minnaard, Smeets, & Lesscher, 2017). For example, Rossiter et al. (2012) showed that a non-dependent/non-clinical sample with harmful levels of alcohol consumption demonstrated reduced sensitivity to punishment in a monetary incentive go/no-go task. One might expect that the insensitivity to adverse consequences often observed in alcohol misuse would be reflected in reduced learning from negative prediction errors in our study, and/or increased learning from positive prediction errors. It may be the case that learning from punishers is affected by delays in punishment, as punishment in animal models (e.g., electric shock) is more immediate than the often delayed punishers in human addiction (e.g., loss of health) (Jean-Richard-dit-Bressel, Ma, Bradfield, Kilcross, & McNally, 2019). It is also possible that the salience of the punisher may affect learning rates (Corr, 2004). In the current study punishment was in the form of negative feedback, and individuals may differ in the extent to which this is encoded as aversive. Another factor which may have implications on learning rates is the ability of subjects to learn the contingencies in the training phase of the PST. Jean-Richard-dit-Bressel et al. (2019) report that the inability to learn contingencies impairs punishment sensitivity in rats, rather than an aversion to negative consequences. I presented subjects with two blocks of the PST training phase, in comparison other studies (Grogan et al., 2017; Kunisato et al., 2012) have trained subjects to criterion, resulting in stronger learning of the reward/punishment contingencies. In our sample, the high alcohol risk group were more accurate in learning CD pairs in the PST training phase, however no group differences were observed for AC or EF pairs. Future research may also control for mood as a modulator of reinforcement learning, as depression has been shown to increase sensitivity to punishments (Eshel & Roiser, 2010) and affect PST performance (Cavanagh et al., 2011; 2019). However, substance dependence is not always characterised by blunted negative learning rates; a recent

study by Kanen et al., (2019) found that SUD patients demonstrated increased punishment learning rates and reduced reward learning rates in comparison to healthy controls on a serial probabilistic reversal learning task, which were sensitive to dopaminergic treatments. Nonetheless, the role of punishment in non-dependence alcohol use warrants further exploration in order to elucidate these contrasting findings.

Repeated correlations between random samples from the computational model trace and the AUDIT showed a negative relationship with positive learning rate, and a positive relationship with negative learning rate. This suggests a potential bidirectional relationship between negative and positive learning rates and alcohol risk (i.e., high risk is associated with greater negative learning rate, and lower positive learning rate). However, more robust correlations were observed between the AUDIT and boundary separation, drift-rate, and non-decision time. Lower boundary separation values were associated with higher AUDIT risk scores, indicating an emphasis on speed relative to accuracy in high risk alcohol users, with lower boundary separation requiring less evidence before a response boundary is reached. Higher drift-rates were associated with high alcohol risk, indicating faster and stronger accumulation of evidence for low conflict trial-types. A negative relationship was observed between non-decision time and the AUDIT, indicating that smaller non-decision times are associated with increased alcohol misuse risk.

Although no published studies to date have examined PST performance using drift-diffusion models to predict alcohol risk, a small number of studies have applied DDMs to value-based decision-making tasks in substance dependent or gambling disorder samples. Comparing gambling disorder patients with controls, Wiehler and Peters (2020) applied RLDDMs to a four armed bandit task, and found that patients had lower learning rates (for both dual and singular models), smaller non-decision times, and a more rapid reduction in the boundary separation parameter across time. Our findings similarly showed a negative relationship between non-

decision times, boundary separation, and alcohol risk, with smaller non-decision times and narrower decision bounds associated with increased alcohol risk. A study by Mandali, Weidacker, Kim, and Voon (2019) sought to examine uncertainty and conflict in a drift-diffusion model applied to the two-step task, showing that alcohol-dependent participants had slower drift-rates relative to healthy controls under conditions of low conflict. Our study also found that slower drift-rates predicted higher alcohol risk, however this was the case for both high (Win-Win) and low (Win-Lose) conditions in the PST test phase, whereas faster drift-rates for Lose-Lose conditions was associated with low-risk. It has been suggested that efficiency of evidence accumulation (i.e., drift-rate) is a trait-like and task-general marker of psychopathology (Weigard & Sripada, 2021), including externalising disorders such as substance dependence. Weigard et al. (2021) found that reduced evidence accumulation predicted substance-use (a combination of alcohol volume, marijuana, and cigarette consumption frequency) in a young adult sample, however our study found that *increased* drift-rates predict higher alcohol misuse risk in a cross-sectional sample. These discrepancies in findings may be due to task-specific differences (i.e., the Go/No-Go and PST assess different aspects of decision-making: inhibition versus reward learning), or differences in type of substance-use measured (reduced evidence accumulation may be more relevant to combined substance use, rather than hazardous alcohol use in isolation) that influence estimates of DDM parameters. Further research may elucidate substance-specific effects on evidence accumulation as measured by a variety of experimental tasks.

In a study examining the utility of behavioural task outcomes vs. self-report surveys to predict 'real-world' outcomes involving self-regulation, Eisenberg et al. (2019) found that factors derived from task performance were poorer predictors of variation in binge and problematic drinking compared with self-report surveys. The task factors reported in this study were derived from a range of behavioural and computational indices (incl. DDM parameters such as boundary separation) from 37 behavioural tasks (including the PST), and the survey

factors included 22 surveys (including the BIS and UPP-S). Similarly, Hedge et al. (2020) have reported no relationship between self-reported impulsivity and boundary separation derived from the DDM. As Wennerhold and Friese (2020) note in the context of cognitive control and inhibition, lack of correspondence between self-report and task-based measures may be due to the difference between typical and maximum performance. Typical performance refers to the tendency of consistent behavior across time, whereas maximum performance suggests the ability of performing maximally at a particular testing session (Wennerhold & Friese, 2020). Although Eisenberg et al. (2019) found no relationship between task performance factors and real-world outcomes, the current study found that such computational parameters were indeed predictive of alcohol misuse. It is possible that the PST places less emphasis on maximum performance than other forced-choice experimental tasks with time-constrained responses such as flanker and stroop tasks.

Increased impulsivity was associated with higher alcohol misuse risk. An extensive literature has linked alcohol use with impulsivity (Dick et al., 2010; Moreno Padilla et al., 2017), suggesting that alcohol consumption leads to impulsive behaviour and also predicts prospective AUD diagnosis (Sher, Bartholow, & Wood, 2000). However certain facets of impulsivity seem to predict different aspects of alcohol consumption. Caswell, Celio, Morgan, and Duka (2015) found that the motor and non-planning subscales of the BIS characterised individuals with high weakly alcohol consumption, whereas O'Halloran et al. (2018) found that the attentional and non-planning subscales predicted alcohol intoxication frequency, but not consumption frequency. Unlike Caswell et al. and O'Halloran et al., who used the AUQ and ESPAD respectively, I used the AUDIT to measure alcohol use. This questionnaire measures aspects of intoxication and consequences of alcohol use, which is more similar to the composite alcohol score in O'Halloran et al. This lends support to the finding that attentional impulsivity is associated with intoxication, rather than level of alcohol consumption alone.

Gender significantly predicted alcohol misuse risk, female group membership was associated with reduced odds of being in the high alcohol group. This is consistent with reports that females consume less alcohol on average compared to men, with fewer heavy drinking episodes and lower lifetime abstinence rates (54.6% versus 34.5%; WHO, 2018). Our finding that age increased the odds of alcohol misuse risk is also consistent with reports that adolescence and young adulthood is associated with increased heavy episodic drinking (WHO, 2018).

In summary, this study found evidence of decision-making differences among individuals with high alcohol misuse compared to low-risk individuals. Specifically, in a probabilistic value-based decision-making task, higher negative learning rate, higher evidence accumulation (drift-rate), lower response threshold (boundary separation), and smaller non-decision times predicted high alcohol-misuse risk. Future studies may consider if these differences generalise to other types of value-based decision-making tasks (e.g., probabilistic reversal learning) and clinical manifestations of alcohol dependence, in addition to examining the pathophysiology of these processes using neural data.

Chapter 4: Reward and punishment processing in former opioid addiction

4.1. Introduction

The World Health Organisation (2018) estimates that over 70% of drug-related deaths are attributable to opioids, and in the United States, opioid-related deaths exceed those caused by car accidents (Bickel et al., 2018; Mattson, 2021; Warner et al., 2011). Increases in the availability of medically-prescribed opioid analgesics in recent years has led to increased accidental overdoses, and increases in the likelihood of transition to illicit opioid dependence (Cicero et al., 2014; Seth et al., 2018; Smyth et al., 2010). Opioid-dependence has been noted for its difficulty to overcome, with one study reporting that 91% of participants relapse following treatment discharge (Smyth et al., 2010).

A substantial literature has identified decision-making impairments relevant to prolonged opioid-use, including greater discounting of delayed rewards, increased risk-tolerance, and differential learning from rewards versus punishments (Madden et al., 1997; Myers et al., 2017; Prosser et al., 2006). For example, risky decision-making in opiate-use is often measured using gambling tasks such as the Iowa Gambling Task (IGT), Cambridge Gambling Task (CGT), Balloon Analogue Risk Task (BART), or Game of Dice task. These experimental tasks manipulate conditions of ambiguity and volatility, so that the outcomes of response selections are uncertain and unpredictable. Relative to healthy controls, opiate users show a preference for smaller short-term rewards over larger long-term rewards, attentional biases towards drug-related cues, and attend less to environmental contingencies in decisions that involve risk-taking (Biernacki et al., 2016; Garland & Howard, 2014; Saleme et al., 2018). These alterations to risky decision-making are also associated with other forms of substance dependence (Chen et al., 2020), and efforts to parse the relative influence of different drug-types are sometimes compounded by evident poly-drug use among research participants. However some studies have shown differences in decision-making between these groups.

Ahn and Vassileva (2016) identified unique features classifying former heroin versus amphetamine users – past heroin-use was predicted by older age, lower education and IQ,

higher non-planning impulsivity, lower attentional and motor impulsivity (BIS), higher urgency and lack of perseverance (UPPS), impaired decision-making on the IGT, and a range of other neurocognitive and personality/psychiatric features. Notably, these features were not predictive of past amphetamine-use (excluding age), and any common predictors between both substance types were in the opposite direction (e.g., younger age classified amphetamine-use).

A number of studies have highlighted differences in decision-making between opiate-users and healthy controls using computational models of decision-making (such as reinforcement learning and drift-diffusion models). Myers et al. (2016) found no mean behavioural differences in a reward and punishment-learning task between groups, however actor-critic computational models showed a reduced recency bias among opioid-users – indicating reduced likelihood of repeating prior responses. Drug-induced changes to the mu-opioid receptor system via acute morphine versus naltrexone administration in healthy humans have also shown changes in drift-rate and starting-point bias in value-based decision-making (Eikemo, Biele, Willoch, Thomsen, & Leknes, 2017). A bidirectional effect was observed, with opioid agonists increasing drift-rate compared to baseline, and blockade decreasing it-boundary separation was comparably unaffected by drug condition (Eikemo et al., 2017). In a computational examination of performance on the IGT between heroin-dependent, amphetamine-dependent, and healthy control groups (drug-use groups in protracted abstinence), Ahn et al. (2014) demonstrated distinct learning deficits between groups. Relative to controls, heroin-users showed reduced loss-aversion, whereas amphetamine-users showed increased reward sensitivity (Ahn et al., 2014). These studies confirm the utility of such computational models in identifying deviances from healthy decision-making in specific types substance dependence (see Gueguen et al., 2021 for recent review).

The extent to which deficits in decision-making persist following successful abstinence is unclear (Biernacki et al., 2016; Koob, 2020). In a review of the literature, Biernacki et al. (2016)

found no evidence for improvements in decision-making tasks among ex-opiate users, nor a relationship between length of abstinence and magnitude of decision-making deficits. This seems to be supported by (Eikemo et al., 2019), who found no evidence for computational or behavioural changes in reward sensitivity in a sample formerly addicted to heroin compared with healthy controls. However, in a prospective longitudinal study by Konova et al. (2020), computationally-derived ambiguity tolerance predicted prospective opiate use, performing similarly with measures of treatment adherence and craving. Of note, Konova et al. (2020) found that baseline ambiguity tolerance was similar between opiate use disorder patients and controls, however the trajectory of this latent cognitive process differed longitudinally, with higher tolerance over time predicting future relapse. This relationship appears to be specific to substance dependence, as no relationship was found between ambiguity tolerance schizophrenia patients. The computational parameter of *risk*-tolerance showed no predictive relationship with future drug-use (Konova et al., 2020). Therefore, this finding suggests that computational models of decision-making may be used to identify high-risk states for opioid relapse across treatment trajectories that may not be present at baseline.

It has been reported that over 50% of individuals who have completed maintenance pharmacotherapy relapse to illicit opioid-use following one-month of treatment (Bentzley et al., 2015), however there has been limited research investigating neurocognitive changes as a consequence of different therapeutic interventions and prolonged opioid abstinence (Stewart et al., 2019). Therefore, probing cognitive changes associated with short versus longer term abstinence presents an opportunity to identify protective factors relevant to developing successful treatment strategies. The present study aimed to examine changes in value-based decision-making among a formerly opioid-dependent sample currently in MMT, the key aims were: (i) to compare decision-making performance in short- versus long-term abstinent groups, and (ii) to identify and compare behavioural, personality, and computational predictors of length of abstinence.

4. 2. Methods

4. 2. 1. Sample

Participants were recruited from a community-based sample receiving methadone maintenance therapy treatment for former opiate-dependence. 81 participants completed the PST and self-report questionnaires. Participants were not provided with a financial incentive for taking part, or reimbursement of travel costs. All participants provided written informed consent to participate, and the study was approved by the Trinity College Dublin School of Medicine Ethics Committee.

4. 2. 2. Materials

4. 2. 2. 1. Probabilistic Selection Task

As described in Chapter 2.1.2.

4. 2. 2. 2. Barratt Impulsiveness Scale (BIS)

As described in Chapter 2. 2.

4. 2. 2. 3. Short- Impulsive Behavior Scale (UPPS-P)

The UPPS-P model measures trait impulsivity with five sub-scales: Negative Urgency, Lack of Premeditation, Lack of Perseverance, Sensation Seeking, and Positive Urgency. The short version of the scale comprises of 20 items across the five subscales scored on a four-point Likert scale from (1; Agree Strongly) to (4; Disagree Strongly) (Cyders et al., 2014).

4. 2. 2. 4. Quality Control

Trials with RTs <150ms and >4000ms were removed from the single trial analysis (similar to Cavanagh et al., 2014). Subjects with over 40 trials removed due to bad RTs were excluded from the analysis (i.e., subject 3 in the PSTNFB, and subject 7 in the PSTWFB – note that subject 7 did not complete the PSTNFB). One subject (subject 80) was excluded from the PSTWFB analysis as they only completed 20 trials of the task. This resulted in a final sample of 81 for the PSTWFB, and 48 for the PSTNFB.

4. 2. 2. 5. Computational Modelling

As described in Chapter 2, Section 2. 4.

4. 2. 2. 6. Machine Learning

As described in Chapter 2, Section 2. 5. Two sets of machine learning models were conducted; (i) linear regression to predict length of opioid abstinence and (ii) logistic regression to predict short-term versus long-term abstinence.

Correlations were conducted to explore the relationship between model parameters, length of abstinence, and the SUPP-S and BIS scores. Parameter values for each subject were sampled from a random chain in the model trace and correlated with each of the questionnaire variables using spearman's r . This step was repeated 1,000 times and the distribution of rho and p-values for each of the parameters and questionnaire variables across all the iterations are presented below. Taking this 'plausible values' approach, I sought to account for variance in the posterior distribution of participant-level parameters (Boehm et al., 2018; Ly et al., 2017).

4. 3. Results

4. 3. 1. Behavioural Results

Participants who did not pass the training phase of the PST ($n=33$) were excluded from subsequent analysis. The final sample consisted of 48 participants (25 M, 23 F) with a mean age of 36.5 (SD = 7.7). Length of time abstinent from methadone ranged from 3 to 7,920 days, with a mean abstinence of 730.15 (1476.77) days. Demographic information and clinical characteristics of the study participants are presented in Table 1, and summary decision-making task performance in Table 2.

For the logistic regression, the sample was divided into short-term versus long-term abstinent participants based on a cut off of 200 days. This resulted in 23 short-term and 25 long-term abstinent participants.

The short-term abstinence group showed significantly higher scores on the AUDIT and DUDIT questionnaires, indicating higher levels of alcohol and drug-use compared to the long-term abstinence group. No group differences were observed on the FTND or impulsivity

questionnaires, indicating that short and long-term abstinence groups showed similar levels of nicotine dependency and self-reported impulsivity.

Table 4. 1. Mean (SD) of sample characteristics and personality questionnaires with Bayesian Mann Whitney U between short versus long-term abstinence groups.

	Total (M, SD)	Short-term (n=23)	Long-term (n=25)	Sig. test (Mann Whitney U)	BF ₁₀
Gender (M/F)	25/23				
Age	36.46 (7.65)	33.52 (7.86)	39.16 (6.49)	412.5	3.71
Age first drug-use	12.61 (2.68)				
Length abstinent (days)	730.15 (1476.77)	94.07 (60.58)	1315.32 (1956.95)	575.0	1518.16
Longest time drug-free (months)	31.55 (50.05)				
No. lapses in past 2-yrs (single-use)	Md = 0, IQR = 1				
No. lapses in past 2-yrs (daily-use)	Md = 0, IQR = 1				
AUDIT	9.21 (11.48)	14.87 (14.18)	4.0 (3.84)	159.5	4.94
DUDIT	27.79 (16.25)	38.74 (6.31)	17.72 (16.12)	77.0	166.34
FTND	5.0 (3.25)	5.0 (3.0)	6.0 (3.0)	336.0	0.83
BIS					
Attention	17.31 (4.24)	18.81 (5.13)	17.16 (4.18)	311.0	0.29
BIS Motor	25.79 (5.13)	27.47 (5.37)	24.16 (5.16)	393.5	1.25

<i>BIS Non-planning</i>	27.68 (5.82)	29.56 (5.87)	26.45 (4.95)	331.0	0.38
SUPP-S					
<i>Negative Urgency</i>	12.17 (2.75)	12.43 (2.33)	11.92 (3.12)	265.5	0.33
<i>Lack of perseverance</i>	6.81 (2.52)	7.0 (2.26)	6.64 (2.78)	249.5	0.32
<i>Lack of premeditation</i>	8.19 (2.46)	8.35 (2.35)	8.04 (2.59)	273.5	0.32
<i>Sensation-seeking</i>	11.35 (2.89)	12.35 (2.67)	10.44 (2.83)	178.0	1.81
<i>Positive Urgency</i>	10.17 (2.93)	10.69 (3.07)	9.68 (2.77)	224.0	0.42

Table 4. 2. Mean (SD) choice and reaction time summary scores for the PST and DD tasks.

	<i>Total (M, SD)</i>	<i>Short-term (n=23)</i>	<i>Long-term (n=25)</i>	<i>Sig. test (Mann Whitney U)</i>	<i>BF₁₀</i>
PST					
<i>Training: AB accuracy</i>	0.79 (0.10)	0.70 (0.16)	0.81 (0.1)	331.0	0.39
<i>Training: CD accuracy</i>	0.72 (0.11)	0.70 (0.09)	0.75 (0.12)	351.0	0.47
<i>Training: EF accuracy</i>	0.65 (0.11)	0.649 (0.10)	0.647 (0.13)	269.0	0.31
<i>Training: AB reaction time</i>	1.28 (0.5)	1.19 (0.53)	1.37 (0.46)	355.5	0.68
<i>Training: CD reaction time</i>	1.39 (0.56)	1.26 (0.64)	1.51 (0.47)	396.0	1.79
<i>Training: EF reaction time</i>	1.42 (0.61)	1.20 (0.53)	1.63 (0.61)	417.0	6.23
<i>Test: Approach AC</i>	0.58 (0.22)	0.62 (0.17)	0.55 (0.16)	336.0	0.57

<i>Test:</i>	0.37 (0.13)	0.35 (0.15)	0.38 (0.14)	275.0	0.35
<i>Approach BD</i>					

4.3.2. Computational Modelling Results

The model showed good convergence using the Gelman-Rubin method (all values <1.1, highest value=1.023). The model DIC was 46890.97.

Table 4. 3. Summary of computational models fit to the training and test phase of the PST.

<i>Model Type</i>	<i>DIC</i>	<i>Samples</i>	<i>Trials</i>	<i>Learning Rate</i>	<i>Parameters</i>	<i>Max. rhat</i>
<i>RLHDDM</i>						
<i>Model 1</i>	46890.97	10,000 (1,000)	All	Dual	a, v, t	1.02
<i>Model 2</i>	62841.65	12,000 (2,000)	All	Single	a, v, t	1.03
<i>Model 3</i>						
<i>HDDM</i>						
<i>Model 1</i>	16544.36	10,000 (1,000)	All	--	a, t, v (WW, WL, LL)	1.02

Mean parameter values for the RLHDDM and HDDM models are presented in the violin plots and table below. The groups did not differ significantly, however the long-term abstinence group showed higher boundary separation in the RLHDDM compared with the short-term group.

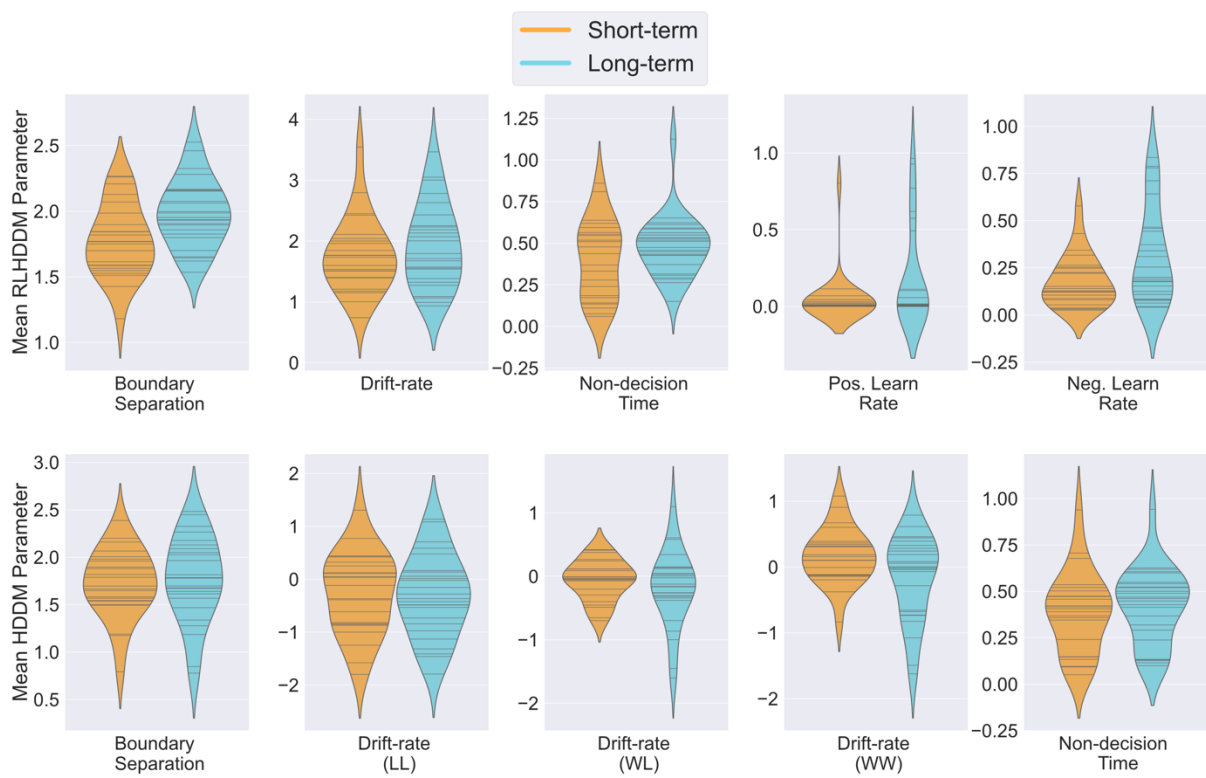


Figure 4. 1. Violin plots of mean RLHDDM and HDDM parameter estimates for short-term versus long-term abstinance groups.

Table 4. 4. Mean (SD) of computational model parameters with significance test between groups.

<i>Parameter</i>	<i>Whole Sample (N = 48)</i>	<i>Short Abs. (n = 23)</i>	<i>Long Abs. (n = 25)</i>	<i>Mann Whitney U (W)</i>	<i>BF₁₀</i>
RLHDDM					
<i>Pos. learning</i>	0.25 (0.22)	0.18 (0.14)	0.32 (0.26)	197.0	1.48
<i>Neg. learning</i>	0.12 (0.27)	0.05 (0.17)	0.19 (0.32)	230.0	0.55
<i>Boundary Separation (a)</i>	1.89 (0.29)	1.77 (0.28)	1.98 (0.26)	162.0	4.14
<i>Non-decision Time (t)</i>	0.44 (0.21)	0.40 (0.24)	0.48 (0.19)	235.0	0.57
<i>Drift-rate (v)</i>	1.85 (0.66)	1.77 (0.62)	1.93 (0.69)	249.0	0.34
HDDM					
<i>a</i>	1.74 (0.41)	1.76 (0.45)	1.72 (0.36)	314.0	0.33
<i>vLL</i>	-0.29 (0.77)	-0.32 (0.78)	-0.27 (0.77)	274.0	0.28
<i>vWL</i>	-0.13 (0.48)	-0.19 (0.60)	-0.06 (0.32)	255.0	0.38
<i>vWW</i>	-0.01 (0.57)	-0.17 (0.64)	0.16 (0.42)	219.0	0.77
<i>t</i>	0.41 (0.21)	0.42 (0.20)	0.39 (0.22)	329.0	0.39

4. 3. 3. Posterior Predictive Checks

Posterior predictive checks were conducted by simulating task data from the posterior of the fitted RLHDDM model and comparing this to the observed data. A random sample from the model trace was used to generate accuracy and RT data for each of the training phase stimulus pairs (AB, CD, EF), this step was repeated 50 times to better capture variability in the posterior distribution. The mean RTs and choices from the simulated data were then compared with the observed data. Figure 4.2 shows the observed versus simulated RT (averaged over 50 iterations) data for each stimulus type (AB, CD, EF). The simulated data overpredicted the correct response RTs, most prominently in the AB pair, and underpredicted error RTs. However, the simulated data fit the pattern of observed data overall, particularly in CD and EF pairs.

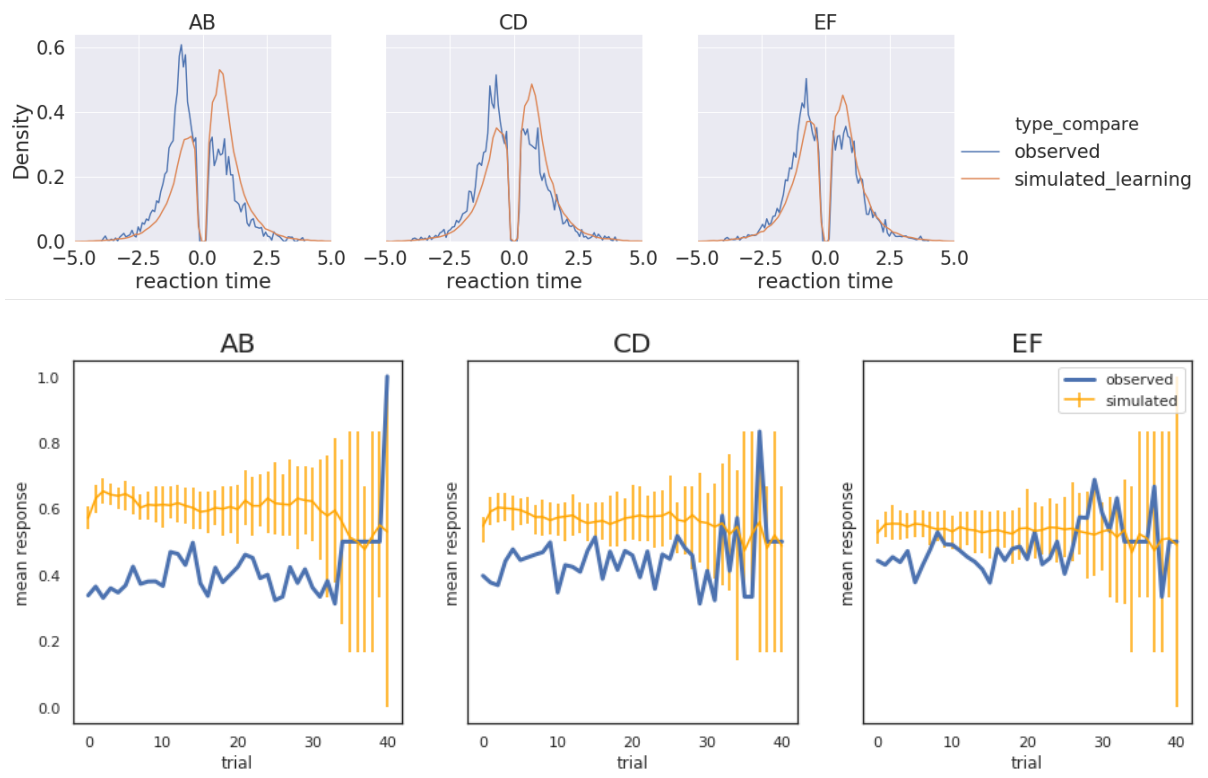


Figure 4. 2. Shows the simulated and observed choices for a subset of trials in the training phase, 1=correct selection, 0=incorrect selection. The simulated data predicted higher correct choice selections compared with the observed data, this was most pronounced in the AB stimulus pair.

Parameter recovery was conducted by fitting an RLDDM model to data simulated from the mean parameters of the fitted model. The mean parameters from the simulated data RLDDM model significantly correlated with those estimated from the observed data RLDDM model (a: $\rho=0.95$, $p<.001$; t: $\rho=0.99$, $p<.001$; v: $\rho=0.61$, $p<.001$, negative learning rate: $\rho=0.47$, $p<.001$; positive learning rate: $\rho=0.69$, $p=.001$). Overall, the parameters showed good recovery, showing that simulated data generated from our RLDDM model with known parameters can be fit to recover these parameters. Scatterplots of the relationship between observed and simulated parameters are presented below.

Simulated vs Observed RLDDM Parameter Correlations

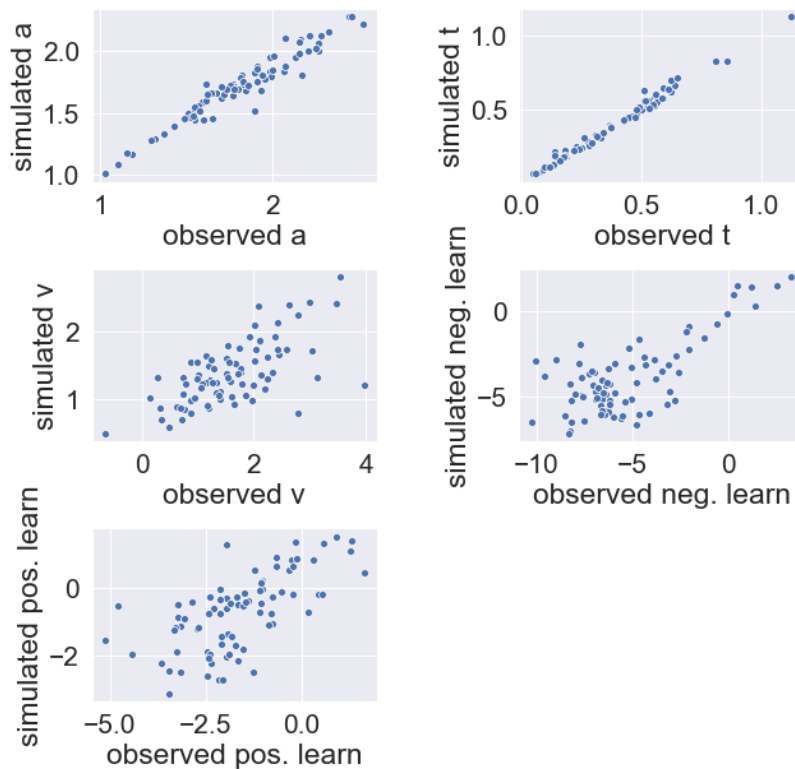


Figure 4. 4. Scatterplots of mean parameter values from RLDDM models with observed vs. simulated data.

The HDDM model fit to the PST test phase showed good convergence using the Gelman-Rubin method (all values <1.1 , highest value=1.024). Posterior predictive checks were conducted by simulating data using the estimated parameters from the HDDM model. 500 samples were drawn in the simulated data, and the summary statistics showed that the observed data fell within the 95% credible interval of the simulated data.

4. 3. 4. Machine Learning

4. 3. 4. 1. Predicting Length of Abstinence

The PST model outperformed the demographic model on 79% of cases based on r score, and the mean r for the test model (0.31, SD = 0.09) was significantly higher than the demographic model (0.15, SD = 0.11); $t(198) = 8.96$, $p < 0.001$. Mean MAE for the test model (0.49, SD = .02) was significantly lower than the demographic model (0.52, SD = .05); $t(198) = -5.19$, $p < .0001$.

Both Approach A and Approach B, in addition to Age were significant predictors, with selection frequencies outside the 95th percentile of demographic model feature selections.

Table 4. 5. PST Summary Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean Choice Frequency - Test</i>	<i>Beta Weight - Test</i>	<i>Mean Choice Frequency - Null</i>	<i>Beta Weight - Null</i>
<i>intercept</i>	5	2.319	5	2.319
<i>Age</i>	4.77	0.140	4.58	0.126
<i>Approach AC</i>	4.53	-0.115	1.43	-0.002
<i>Gender</i>	1.44	-0.013	1.25	-0.013
<i>Approach BD</i>	0.96	-0.002	1.65	0.005

The Computational Model outperformed the demographic model on 75% of cases, with the mean *r* value for the test model (0.24, SD = 0.06) significantly higher than the null model (0.063, SD = 0.12); $t(198) = 9.44, p < .001$. The mean test model MAE (0.49, SD = .02) was significantly lower than the demographic model (0.51, SD = .02); $t(198) = -6.83, p < .0001$. The Drift-rate for Win-Win, Win-Lose, and Lose-Lose conditions, non-decision time, boundary separation, and positive learning rate from the RLHDDM, and age were significant features.

Table 4. 6. Computational Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean Choice Frequency - Test</i>	<i>Beta Weight - Test</i>	<i>Mean Choice Frequency - Null</i>	<i>Beta Weight - Null</i>
<i>intercept</i>	5	2.319	5	2.139
<i>HDDM: Drift-rate Win-Win</i>	4.31	-0.084	0.84	0.0001
<i>Age</i>	4.18	0.060	4.24	0.091
<i>RLHDDM: Non-decision Time</i>	3.72	0.055	0.77	0.002
<i>RLHDDM: Boundary Separation</i>	3.3	0.034	0.76	-0.0002
<i>RLHDDM: Pos. Learning Rate</i>	2.52	0.024	1.0	-0.003
<i>HDDM: Drift-rate Win-Lose</i>	1.98	-0.010	0.74	-0.001
<i>HDDM: Drift-rate Lose-Lose</i>	1.75	-0.012	0.81	0.0002
<i>Gender</i>	0.61	-0.003	0.53	-0.004
<i>RLHDDM: Drift-rate</i>	0.58	0.001	0.83	0.0003
<i>HDDM: Non-decision Time</i>	0.39	-0.001	0.72	0.0003
<i>RLHDDM: Neg. Learning Rate</i>	0.27	0.001	0.77	-0.0007
<i>HDDM: Boundary Separation</i>	0.22	-0.0004	0.69	0.001

The Personality Model r underperformed the demographic model in 61% of cases, with the test model r (0.026, 0.1) significantly lower than the demographic model (0.096, 0.12); $t(198) = -4.21$, $p < .0001$. The mean MAE for the test model (0.52, SD = .02) was significantly lower for the demographic model (0.51, SD = .02) compared with the test model (0.51, SD = .02); $t(198) = 3.99$, $p < .0001$. Age, the SUPP-S sensation-seeking and positive urgency scales, BIS motor impulsivity were significant features in the model.

Table 4. 7. Personality Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean Choice Frequency - Test</i>	<i>Beta Weight - Test</i>	<i>Mean Choice Frequency - Null</i>	<i>Beta Weight - Null</i>
<i>intercept</i>	5	2.319	5	2.319
<i>Age</i>	3.96	0.073	4.2	0.098
<i>SUPPS: Sensation-seeking</i>	2.75	-0.033	0.95	0.002
<i>SUPPS: Pos. Urgency</i>	1.9	-0.014	1.08	0.002
<i>BIS: Motor</i>	1.12	-0.007	0.84	0.003
<i>BIS: Attention</i>	0.55	0.005	0.87	-0.002
<i>SUPPS: Neg. Urgency</i>	0.54	-0.003	0.86	-0.003
<i>Gender</i>	0.29	-0.0009	0.75	-0.006
<i>SUPPS: Lack of Perseveration</i>	0.19	-0.0004	0.89	-0.004
<i>SUPPS: Lack of Premeditation</i>	0.16	0.0005	0.83	0.0003
<i>BIS: Non-planning</i>	0.15	0.0005	0.71	-0.0006

The Combined Model r (0.15, SD = 0.13) outperformed the demographic model (0.03, SD = 0.14) on 69% of iterations, and was significantly higher; $t(198) = 5.42, p < .0001$. The mean MAE for the test model (0.506, SD = .02) was significantly lower than the demographic model (0.512, SD = .02); $t(198) = -2.16, p = .032$. Drift-rate for Win-Win, Win-Lose, and Lose-Lose conditions, age, non-decision time (RLHDDM), boundary separation (RLHDDM), positive learning rate (RLHDDM), positive urgency (SUPP-S), sensation-seeking (SUPP-S), and motor impulsivity (BIS) were significant features.

Table 4. 8. Combined Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean Choice Frequency - Test</i>	<i>Beta Weight - Test</i>	<i>Mean Choice Frequency - Null</i>	<i>Beta Weight - Null</i>
<i>intercept</i>	5	2.319	5	2.319
<i>HDDM: Drift-rate Win-Win</i>	3.95	-0.068	0.59	0.0001
<i>Age</i>	3.65	0.044	3.96	0.077
<i>RLHDDM: Non-decision Time</i>	3.32	0.044	0.67	-0.001
<i>RLHDDM: Boundary Separation</i>	2.77	0.025	0.7	0.002
<i>SUPPS: Sensation-seeking</i>	2.36	-0.021	0.65	0.0005
<i>RLHDDM: Pos. Learning Rate</i>	1.81	0.013	0.79	-0.0003
<i>SUPPS: Pos. Urgency</i>	1.53	-0.008	0.63	-0.0006
<i>HDDM: Drift-rate Win-Loss</i>	1.42	-0.007	0.49	0.001
<i>HDDM: Drift-rate Loss-Loss</i>	1.12	-0.006	0.67	-0.001
<i>BIS: Motor</i>	0.87	-0.004	0.52	-0.001
<i>BIS: Attention</i>	0.45	0.003	0.59	-0.003
<i>RLHDDM: Drift-rate</i>	0.35	0.001	0.7	0.0003
<i>Gender</i>	0.32	-0.002	0.39	-0.003
<i>RLHDDM: Neg. Learning Rate</i>	0.27	0.001	0.73	-0.001
<i>BIS: Non-planning</i>	0.25	0.001	0.59	-0.001
<i>HDDM: Non-decision Time</i>	0.19	-0.0008	0.6	0.002
<i>SUPPS: Neg. Urgency</i>	0.18	-0.0004	0.43	-0.003
<i>SUPPS: Lack of Perseveration</i>	0.17	0.0004	0.62	-0.00002
<i>HDDM: Boundary Separation</i>	0.14	0.0002	0.65	-0.002
<i>SUPPS: Lack of Premeditation</i>	0.07	0.0001	0.41	0.0001

4. 3. 4. 2. Group Classification

Short-term abstinence was the positive class used in logistic machine learning models,

therefore the direction of beta weights may differ from the linear analysis presented above.

The PST Summary Model AUC (0.69, SD = .04) outperformed the demographic model (0.66, SD = 0.05) on 66% of iterations, and was significantly higher; $t(198) = 5.76, p < .0001$. The test model mean Brier score (0.227, SD = .01) was significantly lower than the demographic model (0.236, SD = .01); $t(198) = -4.86, p < .0001$. Age, gender, Approach AC, and Approach BD were all significant features in the model.

Table 4. 9. PST Summary Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean Choice Frequency - Test</i>	<i>Beta Weight - Test</i>	<i>Mean Choice Frequency - Null</i>	<i>Beta Weight - Null</i>
<i>intercept</i>	5	-0.075	5	-0.083
<i>Age</i>	4.99	-0.737	4.96	-0.582
<i>Gender</i>	4.82	0.409	4.38	0.284
<i>Approach AC</i>	4.75	0.276	3.16	-0.002
<i>Approach BD</i>	4.17	-0.259	3.2	-0.022

The Computational Model AUC (0.69, SD = 0.04) was significantly higher than the demographic model (0.60, SD = 0.06) and outperformed in 78% of cases; $t(198) = 10.92, p < .001$. The mean test model Brier score (0.22, SD = 0.02) was significantly lower than the demographic model (0.25, SD = 0.02); $t(198) = -10.13, p < .0001$. Age, gender, all of the RLHDDM parameters (boundary separation, positive learning rate, negative learning rate, drift-rate, non-decision time), drift-rate for the Win-Win condition, and boundary separation from the HDDM were significant predictors.

Table 4. 10. Computational Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean Choice Frequency - Test</i>	<i>Beta Weight - Test</i>	<i>Mean Choice Frequency - Null</i>	<i>Beta Weight - Null</i>
<i>Intercept</i>	5	-0.131	5	-0.088
<i>Age</i>	4.97	-0.448	4.65	-0.401
<i>RLHDDM: Boundary Separation</i>	4.97	-0.505	1.76	0.002
<i>RLHDDM: Pos. Learning Rate</i>	4.86	-0.387	1.97	0.002
<i>HDDM: Drift-rate (Win-Win)</i>	4.85	0.346	1.68	0.003
<i>RLHDDM: Neg. Learning Rate</i>	4.71	-0.265	2.06	-0.019
<i>Gender</i>	4.45	0.280	3.09	0.151
<i>HDDM: Boundary Separation</i>	3.22	0.157	1.82	0.001
<i>RLHDDM: Drift-rate</i>	2.22	-0.035	1.86	-0.002
<i>RLHDDM: Non-decision Time</i>	2.09	-0.031	1.69	-0.008
<i>HDDM: Drift-rate (Lose-Lose)</i>	1.7	0.002	1.6	0.003
<i>HDDM: Non-decision Time</i>	1.62	-0.003	1.9	-0.00003
<i>HDDM: Drift-rate (Win-Lose)</i>	1.4	0.005	1.67	0.0005

The Personality Model AUC (0.638, SD = .05) did not significantly differ from the demographic model (0.632, SD = 0.06), and outperformed in 52% of iterations; $t(198) = 10.95$, $p = 0.42$. The Brier score for the test model (0.248, SD = .01) also did not differ significantly from the demographic model (0.246, SD = .02); $t(198) = 0.72$, $p = 0.47$. Age, gender, sensation-seeking (SUPP-S) and motor impulsivity (BIS) were significant features in the model.

Table 4. 11. Personality Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean Choice Frequency - Test</i>	<i>Beta Weight - Test</i>	<i>Mean Choice Frequency - Null</i>	<i>Beta Weight - Null</i>
<i>intercept</i>	5	-0.086	5	-0.087
<i>Age</i>	4.76	-0.333	4.79	-0.471
<i>SUPPS: Sensation-seeking</i>	4.57	0.236	2.34	0.011
<i>BIS: motor</i>	3.99	0.170	2.15	-0.011
<i>Gender</i>	3.31	0.114	3.7	0.201
<i>SUPPS: Lack of Premeditation</i>	1.64	-0.043	1.85	-0.0002
<i>BIS: Attention</i>	1.58	-0.035	2.14	0.026
<i>SUPPS: Neg. Urgency</i>	1.49	-0.037	2.31	-0.002
<i>SUPPS: Pos. Urgency</i>	1.38	0.021	2.17	-0.015
<i>SUPPS: Lack of Perseveration</i>	1.29	0.020	2.01	-0.008
<i>BIS: Non-planning</i>	1.04	0.010	2.01	0.001

The Combined Model AUC (0.68, SD = .05) was significantly higher than the demographic model (0.59, SDD = 0.07), and outperformed in 75% of iterations; $t(198) = 8.79, p < .0001$. The mean test model Brier score (0.23, SD = .02) was also significantly lower than the demographic model (0.25, SD = .01); $t(198) = -7.75, p < .0001$. Age, gender, boundary separation from the RLHDDM and HDDM, positive and negative learning rates (RLHDDM), drift-rate for the Win-Win condition (HDDM), sensation-seeking and lack of premeditation (SUPPS), and motor impulsivity (BIS) were significant predictors.

Table 4. 12. Combined Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean Choice Frequency - Test</i>	<i>Beta Weight - Test</i>	<i>Mean Choice Frequency - Null</i>	<i>Beta Weight - Null</i>
<i>intercept</i>	5	-0.108	5	-0.091
<i>RLHDDM: Boundary Separation</i>	4.89	-0.330	1.6	0.0001
<i>Age</i>	4.86	-0.275	4.48	-0.342
<i>SUPPS: Sensation-Seeking</i>	4.82	0.271	1.58	-0.011
<i>RLHDDM: Neg. Learning Rate</i>	4.77	-0.269	1.84	-0.004
<i>RLHDDM: Pos. Learning Rate</i>	4.71	-0.251	1.7	-0.004
<i>HDDM: Drift-rate (Win-Win)</i>	4.7	0.253	1.57	-0.008
<i>BIS: Motor</i>	4.17	0.201	1.57	-0.016
<i>Gender</i>	3.73	0.135	2.77	0.110
<i>HDDM: Boundary Separation</i>	2.22	0.065	1.59	-0.012
<i>SUPPS: Lack of Premeditation</i>	1.99	-0.057	1.25	0.0003
<i>SUPPS: Neg. Urgency</i>	1.79	-0.059	1.34	-0.003
<i>SUPPS: Pos. Urgency</i>	1.78	0.031	1.43	0.004
<i>RLHDDM: Non-decision Time</i>	1.74	-0.033	1.51	0.007
<i>RLHDDM: Drift-rate</i>	1.62	-0.018	1.41	0.001
<i>SUPPS: Lack of Perseveration</i>	1.58	0.007	1.71	-0.010
<i>HDDM: Drift-rate (Lose-Lose)</i>	1.37	-0.008	1.6	-0.024
<i>HDDM: Non-decision Time</i>	1.33	0.013	1.39	0.0145
<i>HDDM: Drift-rate (Win-Lose)</i>	1.26	0.014	1.74	-0.008
<i>BIS: Non-planning</i>	1.21	-0.006	1.36	-0.005
<i>BIS: Attention</i>	1.2	-0.013	1.3	-0.0016

Kruskal Wallis tests comparing mean model accuracy between test models showed that MAE and r scores differed significantly between groups ($K = 90.42, p < .0001$; $K = 193.25, p < .0001$), as did Brier and AUC scores for the logistic models ($K = 104.24, p < .0001$; $K = 82.81, p < .0001$).

Post-hoc tests showed that the combined models outperformed the individual models (p 's < .002). The Computational and PST Summary models outperformed the Personality models (p 's < .001), however the Computational and PST Summary models did not differ significantly from each other (p 's > 0.15).




Figure 4. 5. Violin plots summarising machine learning model metrics for PST Summary, Computational, Personality, and Combined Models.

4. 3. 5. Correlations with Length of Abstinence

No strong correlations were observed between the computational parameters from the RLHDDM and length of abstinence, with the positive and negative learning rates, and drift-rate correlation distributions overlapping with the null correlations. More robust correlations were observed between the computational parameters and personality questionnaires. In particular, the boundary separation parameter correlated negatively with the BIS Motor subscale. This parameter also correlated negatively with the positive urgency sub-scale of the SUPP-S, and the non-decision time parameter showed negative correlations with all of the SUPP-S scales.




Figure 4. 6. Distribution of rho (upper row) and log10(p-values; lower row) for correlations between RLHDDM computational model parameters and length of abstinence.




Figure 4. 7. Distribution of rho (upper row) and log10(p-values; lower row) for correlations between HDDM computational model parameters and length of abstinence.

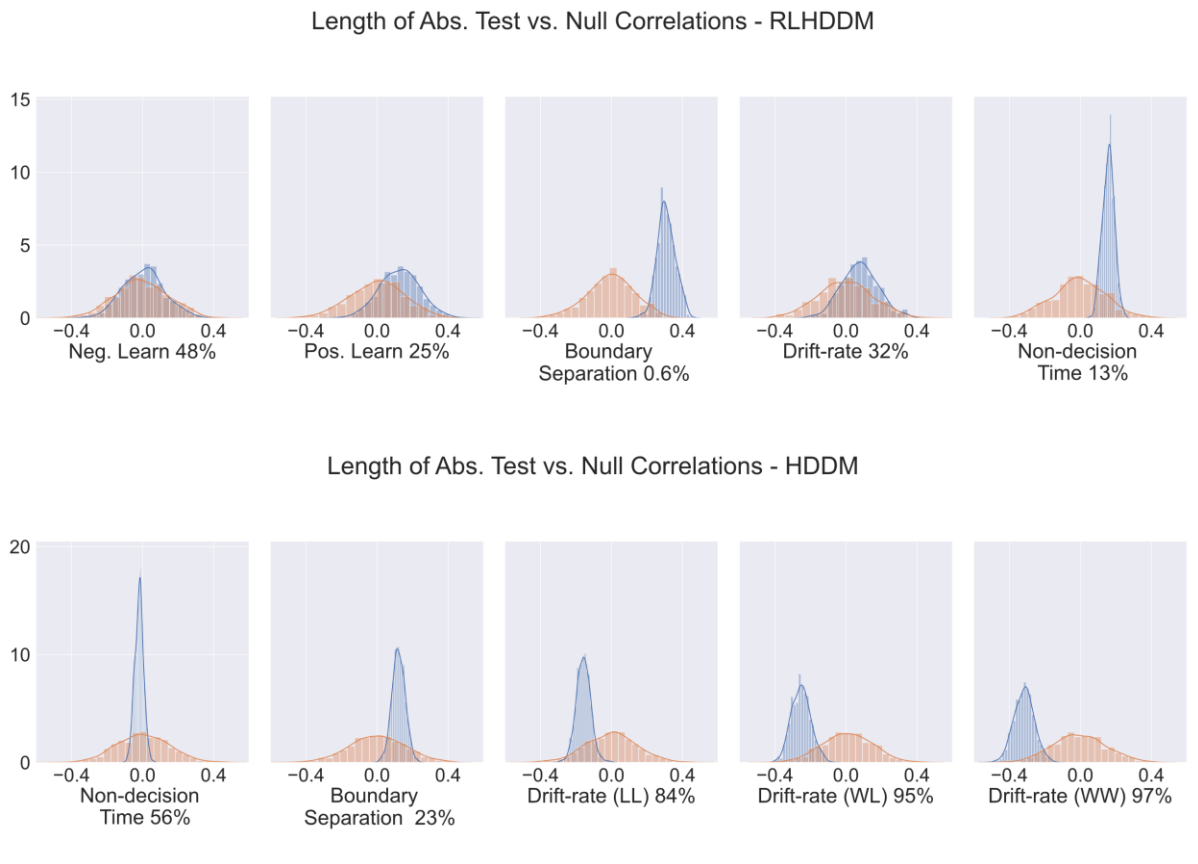


Figure 4. 8. Distribution of test versus null rho values for each RLHDDM parameter correlation with length of abstinence.

4. 3. 6. Correlations between ML Features

Feature	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1. Neg. Urgency	—																			
2. Lack of Persev..	0.076	—																		
3. Lack of Premed.	0.279	0.466***	—																	
4. Sensation-Seeking	0.316*	0.067	0.291*	—																
5. Pos. Urgency	0.421**	0.155	0.413**	0.256	—															
6. BIS Attention	0.425**	0.309*	0.539***	0.141	0.495***	—														
7. BIS Motor	0.336*	0.117	0.434**	0.173	0.376**	0.507***	—													
8. BIS Non-planning	0.258	0.539***	0.720***	0.249	0.271	0.443**	0.403**	—												
9. RLHDDM: a	0.019	-0.009	-0.053	0.064	-0.019	0.062	-0.299*	-0.192	—											
10. RLHDDM: v	0.031	0.115	0.139	-0.177	0.093	0.300	0.009	0.131	-0.092	—										
11. RLHDDM: t	-0.059	-0.051	-0.097	0.015	-0.123	0.174	-0.015	-0.040	0.316*	0.177	—									
12. RLHDDM: Neg. Learning	0.139	-0.095	-0.035	0.108	0.048	0.041	0.111	-0.012	0.001	-0.018	0.984	—								
13. RLHDDM Pos. Learning	0.086	-0.092	0.044	-0.156	-0.098	0.234	-0.025	-0.041	0.033	0.383**	0.110	0.105	—							
14. HDDM a	0.066	0.159	-0.057	-0.081	-0.051	0.132	-0.119	-0.244	0.453**	-0.212	0.496***	-0.009	0.174	—						
15. HDDM vWW	0.285*	0.261	-0.046	0.086	-0.044	0.122	-0.042	0.173	-0.209	0.223	0.070	0.051	-0.019	-0.067	—					
16. HDDM vLL	0.023	0.121	0.028	0.279	0.165	0.038	-0.063	-0.111	0.134	-0.277	0.186	0.188	-0.216	0.175	0.113	—				
17. HDDM vWL	0.222	0.246	0.003	0.111	0.175	0.108	-0.014	-0.047	0.007	-0.036	0.162	0.218	-0.198	0.007	0.586***	0.621***	—			
18. HDDM: t	0.222	-0.219	-0.206	-0.116	-0.034	0.226	-0.058	-0.170	0.035	0.127	0.607***	-0.004	0.040	0.317*	0.134	0.181	0.192	—		
19. Approach AC	0.142	0.306*	0.043	0.084	-0.072	0.103	-0.104	0.214	-0.199	0.219	0.045	-0.058	0.047	-0.064	0.807***	0.199	0.395**	0.100	—	
20. Approach BD	0.003	0.117	-0.206	0.123	0.288*	0.122	0.051	-0.028	0.156	-0.197	0.043	0.274	-0.290*	-0.048	-0.003	0.725***	0.653***	0.065	-0.134	—

Table 4. 13. Partial Spearman's correlations between machine learning features, with age as a covariate. Significant correlations are highlighted between (i) self-report measures (orange), (ii) task-related parameters and variables (blue), (iii) task-related and self-report measures (green). * $p < .05$, ** $p < .01$, *** $p < .0001$

4. 4. Discussion

The current study sought to predict length of abstinence from non-treatment opioid-use (hereafter, simply 'length of abstinence') in a sample receiving methadone maintenance therapy. Utilising task-based and self-report features, four models were compared: a PST Summary Model with mean accuracy from the PST test phase, a Computational Model with mean DDM model parameters derived from the PST training and test phase, a Personality Model with questionnaire measures of impulsivity, and a model combining personality and computational model features. In the Combined Model, drift-rate for Win-Win conditions in the PST test phase was the strongest feature, showing a negative relationship with length of abstinence. Larger non-decision times, wider boundary separation, and higher positive learning rates were also associated with increased length of abstinence. In the PST Summary Models, lower Approach AC accuracy and higher Approach BC accuracy was associated with long-term abstinence.

Smaller drift-rates across all conflict conditions of the PST test phase predicted longer abstinence across the Computational and Combined Models. This suggests that reduced efficiency of evidence accumulation, a risk-factor that has been linked with prospective substance use in young adults (Weigard et al., 2021), predicts longer lengths of abstinence among individuals in MMT. It is somewhat surprising that long-term abstinence, likely conceptually associated with increased self-regulation, shows reductions in drift-rate as observed in other externalising disorders (e.g., ADHD; Ziegler et al., 2016). Although drift-rates seem to generalise across a range of experimental tasks (Schubert et al., 2016), given that this study is the first to examine DDM features derived from PST performance in former opioid-users, it is conceivable that task-specific factors related to the PST test phase may influence the efficiency of evidence accumulation. Further

research, comparing DDM parameters derived from other tasks among former opioid-users may elucidate some of these peculiarities.

More intuitively, larger response boundaries and larger non-decision times- associated with slower and more conservative responding in the PST training phase, also predicted longer length of abstinence. Shared variance between age and these parameters may explain their significance, as age has been shown to robustly affect boundary separation and non-decision time parameters, unlike drift-rate (Dully et al., 2018). However, robust partial correlations with age as a covariate showed positive correlations between boundary separation, non-decision time, and length of abstinence. This finding shows a positive relationship between these parameters and length of opioid abstinence beyond the effect of age, and suggests an increased ability to inhibit responses and accumulate sufficient evidence in longer-term abstinence. Increases in boundary separation and non-decision time parameters derived from the PST have also been observed in ADHD patients receiving stimulant medication (Pederson et al., 2017). This suggests that modulation of noradrenergic function associated with task adaptations in ADHD may be similar in opioid long-term opioid abstinence. Although positive learning rate was a significant feature in the machine learning models, it did not show robust correlations with length of abstinence. Nonetheless, correlations between the mean parameter estimates showed consistency between training and test PST phases, and showed some relationships with self-report measures (see Supplementary Materials for robust correlations between RLHDDM model parameters, the BIS, and the S-UPPS for the whole participant sample). Boundary separation showed a negative relationship with motor impulsivity, and drift-rate for Win-Wins showed a positive relationship with negative urgency.

While the current study found evidence of a predictive relationship between DDM parameters and length of abstinence, mean group differences in parameter estimates were not observed. In a perceptual decision-making task with probabilistic feedback, Eikemo et al. (2019) also

found no group differences in drift-diffusion model parameters between former heroin-users in opioid maintenance therapy and healthy controls. As the authors suggest, cognitive impairments associated with opioid-use may be due to social factors involved in a substance-dependent lifestyle that are no longer present in ex-users (e.g., poor health, exposure to violence; Darke et al., 2000). Therefore, the absence of alterations in task-related decision-making behaviour may be due to lack of exposure to these factors in ex-users. Nonetheless, while no association was found in computational parameters between groups, it is possible that these are still useful predictors of abstinence length ('association is not prediction'; Bzdok et al., 2020).

The Personality Models showed relatively poorer predictions compared with the task-related models, suggesting that self-reported impulsivity does not predict length of abstinence as well as the PST and computational features. A number of studies have shown that self-reported impulsivity is a predictor of opioid abstinence, with lower impulsivity associated with increased likelihood of opioid abstinence ((Su et al., 2015; Zhu et al., 2018a). The influence of impulsivity on abstinence appears stronger in younger individuals, and is correlated with comorbid factors such as depression and ADHD (Evren et al., 2018; Peters & Soyka, 2019; Zhu et al., 2018b). In a study comparing different aspects of task-based impulsivity, Li et al., (2021) found that patients on MMT had greater delay discounting and cue-induced craving compared to healthy controls, which correlated positively with self-reported urgency. No differences in risk-taking were observed and deficits in paired associate learning were accounted for by depression and anxiety symptoms in the patient group (Li et al., 2021). It is possible that task-related impulsivity is more predictive of length of opioid abstinence than questionnaires (Lane et al., 2003). However, in both of the Combined Models, motor impulsivity and sensation-seeking were identified as significant features, showing a negative relationship with length of abstinence- this is consistent with the direction of beta values in models classifying heroin-disorder patients reported in Ahn et al. (2016).

The PST Summary models performed similarly to the Computational Models in the current study, indicating that mean choice accuracy is as predictive of abstinence length compared with DDM model parameters. Several studies have highlighted differences in PST performance among clinical, psychiatric, and substance dependent groups (Baker et al., 2013; Chase et al., 2010; Dowd et al., 2016; Rustemeier et al., 2012; Strauss et al., 2015), and this is the first study to find a relationship with opioid abstinence. Reduced learning from both positive and negative feedback predicted long-term abstinence. Additionally, avoiding punished stimuli from the training phase correlated negatively with self-reported positive urgency, and learning more from rewards correlated positively with lack of perseverance.

The current study possesses a number of strengths (namely, the identification of predictors of abstinence in a rare substance group, and the use of a rigorous machine learning approach with nested cross-validation), nonetheless a number of limitations must also be acknowledged. Firstly, the lack of comparison with (i) an age-matched healthy control group, and (ii) a group of ex-opioid users who have completed MMT, limits the conclusions that can be made regarding decision-making impairments due to MMT, versus former opioid-use, versus normative cognitive functioning. Second, the use of a larger sample with an external validation test dataset would contribute greatly to the generalisability of findings and improve individual MMT patient predictions.

In conclusion, this study provides a first application of DDMs to the PST to predict length of opioid abstinence, highlighting a number of relevant features. Increased rates of evidence accumulation, larger non-decision times, and increased boundary separation were significant predictors, providing unique insights into mechanisms that may be involved in prolonged abstinence from opioid-use. This contributes to a growing literature examining the effects of successful abstinence on cognition and decision-making in opioid dependence.

Chapter 5: Reward and punishment learning predictors of smoking status

5. 1. Introduction

The current study examined reward and punishment learning in ex-smokers, current smokers, vapers, and non-smokers using the PST. Based on previous studies in substance-dependent samples (Baker et al., 2011, Baker et al., 2013; Nestor et al., 2018b), and availability of research participant data, Study 1 (N = 105) compared PST test phase performance between current smokers, ex-smokers, and non-smokers. I hypothesized that the current smokers would show impaired reward learning compared with ex-smokers and non-smokers. I also expected that ex-smokers, relative to current and never smokers, would learn best from punishers. Study 2 extended on the behavioral findings of Study 1 with a larger pool of participants (N = 173) by investigating latent higher-order cognitive traits via the application of computational models to the PST, and used these to predict smoking group status using penalized logistic regression. Study 2 also included an additional 'Vaping' group that reported regular e-cigarette/vape consumption.

5. 2 Study 1

5. 2. 1. Methods

5. 2. 1. 1. Participants

Participants were recruited via posters displayed on the University College Dublin campus, Trinity College Dublin campus, and in the local community. Participants were provided with €10 in compensation for taking part in the study. 57 current smokers had smoked over 40 lifetime cigarettes, with at least weekly smoking in the past 30 days. 40 ex-smokers smoked more than 40 cigarettes in their lifetime, with fewer than one cigarette per week (4 participants), or no cigarettes at all, in the past 30 days. 43 non-smokers smoked on fewer than 40 occasions in their entire lifetime with no cigarettes at all in the past 30 days. Exhaled carbon monoxide readings were collected from a subset of 60 participants (25 Smokers; 17 ex-smokers; 13 non-smokers). Smokers had readings of ≥ 6 ppm, and ex-smokers and non-smokers ≤ 5 ppm (Low, Ong, & Tan, 2004).

5. 2. 1. 2. Procedure

Ethics Committees from University College Dublin and Trinity College Dublin approved the study. Current smokers were requested to smoke as normal prior to the experiment, and therefore were not in acute abstinence. Participants completed the experimental tasks alone in a sound-attenuated booth. Questionnaires were completed during the testing session, or at home via an online survey platform. The PST was part of a larger test battery that took approximately 1 h. Participants were compensated with €10 (approximately \$12) plus maximum travel expenses of €10.

5. 2. 1. 3. Measures

5. 2. 1. 3. 1. *Probabilistic selection task*

As described in Chapter 2.1.

Statistical analyses were performed in IBM SPSS (Version 23). Non-parametric tests were used when appropriate. Alpha was 0.05 unless stated otherwise due to multiple comparison correction. Our goal was to predict group membership (see Yarkoni and Westfall, 2017, for a rationale for prediction versus explanation) and therefore percent Approach A and Avoid B selections were predictor variables in a multinomial logistic regression model. The non-smoker group was the reference category and p values were calculated using 1000 bootstrapped samples.

5. 2. 1. 3. 2. *Questionnaire measures*

5. 2. 1. 3. 2. 1. Fagerstrom Test for Nicotine Dependence (FTND)

As described in Chapter 2, Section 2.3.

5. 2. 1. 3. 2. 2. The European School Survey Project on Alcohol and Other Drugs (ESPAD)

The ESPAD questionnaire on substance use (Hibell & Bjarnason, 2008) was used to assess lifetime and past 30 days smoking, and past 30 days alcohol use. Participants were asked 'On how many occasions (if any) during your lifetime have you smoked cigarettes?', 'On how many occasions (if any) during your lifetime have you vaped (i.e., smoked E-cigarettes)?' (Responses = '0', '1-2', '3-5', '6-9', '10-19', '20-39', or '40+' occasions), and 'On how many of the past 30 days did you smoke tobacco cigarettes, even one or two puffs?', 'On how many of the past 30 days did you smoke E-cigarettes, even one or two puffs?' (Responses = '0 days', '1-10 days', '11-20 days', '21-29 days', or 'Daily').

Current smokers reported smoking on 40 or more occasions in their lifetime, and daily in the past month, with no e-cigarette use in the past month. Ex-smokers reported smoking on 40 or more occasions in their lifetime, and not at all in the past month, with no e-cigarette use in the past month. Vapers reported smoking between 0-40+ cigarettes in their lifetime, and not at all in the past month, with daily e-cigarette use and having smoked 40+ e-cigarettes in their lifetime. Non-smokers reported smoking fewer than 20 cigarettes in their lifetime, with no tobacco or e-cigarette use in the past month.

The Sensation Seeking Scale (SSS; Zuckerman, 1971) is a 40-item measure with four sensation seeking subscales: thrill- and adventure-seeking, disinhibition, experience-seeking and susceptibility to boredom (further details in Supplemental Materials).

5.2.2. Results

The final sample consisted of 41 smokers, 29 ex-smokers and 35 non-smokers. Participant characteristics are presented in Table 1 (see Supplemental Materials for further information on the PST training phase). Specific age was collected for 58 participants (55.2%) of the final sample; remaining participants were aged 18–21 years. There was a significant difference between groups

based on the 58 participants who reported their exact age (Kruskal Wallis test; $\chi^2(2, 58) = 8.069$, $p = .018$). The ex-smoker group ($N = 17$) were older than the smoker ($N = 25$) and non-smoker ($N = 16$) groups. The mean FTND for smokers was 2.17 ($SD = 2.26$), indicating the 'Low' dependence that is typical for younger smokers (Li et al., 2015).

Table 5. 1. ESPAD, CO reading, PST performance, reaction times, and personality scores (Barratt Impulsiveness Scale and Sensation-Seeking Scale) by group. [‡]Median(Inter Quartile Range).

Measure	Smokers	Ex-smokers	Non-smokers	Significant difference
Gender (M/F)	26/15	18/11	13/22	-
Age (Years) [‡]	21 (5)	32 (21)	22.5 (3.5)	Ex>S & NS
Lifetime smoking (ESPAD) [‡]	7 (0)	7 (0)	2 (3)	
Past 30 days smoking (ESPAD) [‡]	4 (1)	1 (0)	1 (0)	
Past 30 days Alcohol (ESPAD) [‡]	4 (3)	4 (3)	4 (2)	-
Carbon Monoxide (ppm) [‡]	11 (5.5)	2 (1)	2 (1.5)	S>Ex & NS
FTND Total [‡]	2.17 (2.26)	-	-	N/A
BIS Total [‡]	70 (16)	62 (11)	66.5 (10)	S>Ex
BIS Attentional [‡]	19 (5.5)	16 (4.5)	17.5 (4)	-
BIS Motor [‡]	24 (6)	22 (4)	23 (5.5)	-
BIS Non-planning [‡]	26 (7)	26 (9)	25 (5.75)	-
SSS Total [‡]	26 (11)	20 (11)	21.5 (5)	S>Ex & NS
SSS Boredom Susceptibility [‡]	3 (3)	3 (2.5)	2 (1.75)	-
SSS Disinhibition [‡]	7 (3)	6 (4)	6 (3)	-
SSS Experience Seeking [‡]	7 (3)	6 (3)	6 (3)	S>NS
SSS Thrill & Adventure [‡] Seeking	8 (5)	6 (7)	8 (3.75)	-

For Approach A trials, non-smokers chose A more often than smokers and non-smokers, with a median % choice (interquartile range) of 89(29), 83(32), and 77(38), respectively. For Avoid B trials, non-smokers performed similarly to smokers, with a median % choice of 75(26) and 76(21)

respectively, while ex-smokers avoided the B stimulus on 72% of trials (IQR = 20). Approach A and Avoid B percentages were entered in a multinomial logistic regression. Performance on the learning from reward (i.e., Approach A) test trials successfully predicted smoker group (Approach A, $\chi^2 = 7.01$, $df(2)$ $p = .030$), while learning from punishment (i.e., Avoid B) test trials was just greater than the significance threshold (Avoid B, $\chi^2 = 5.96$, $df(2)$, $p = .051$). Table 2 displays the classification accuracy of the multinomial logistic regression. As the tendency to learn from positive outcomes increased, the likelihood of being a non-smoker relative to smoker ($p = .024$; 95% Confidence Interval – 0.089 to –0.009) or ex-smoker ($p = .04$; 95% CI -0.075 to –0.001) increased. In contrast, as the tendency to learn from punishment increased, the likelihood of being a smoker relative to non-smoker increased ($p = .034$; 95% CI 0.006 to 0.095), but this was not significant for ex-smokers compared with non-smokers ($p > .05$). A separate multinomial regression was conducted to control for the inclusion of lighter smokers in our analysis, and produced similar findings (see Supplementary Materials). CO readings significantly correlated with the tendency to learn from punishment ($\rho = 0.31$, $p = .020$), but not from reward ($\rho = 0.01$, $p = .93$).

BIS-11 scores were compared using Kruskal Wallis tests. Groups differed in total BIS score ($\chi^2(2, 105) = 7.03$, $p = .03$) and total SSS scores ($\chi^2(2, 98) = 7.5$, $p = .02$). Questionnaire data for eight participants were missing (final sample size: 41 Smokers, 29 Ex-Smokers, 28 Non-smokers). Total BIS and SSS scores were entered as predictor variables in a separate multinomial regression model. Neither questionnaire significantly predicted smoking group status ($p > .05$). However, total SSS score predicted the likelihood of belonging to the smoker group relative to the non-smoker group ($p = .046$, 95% CI -0.008 to 0.176). Correlations between the PST (Approach A and Avoid B), and personality measures (the BIS and SSS) were not significant.

Table 5. 2. Classification table for multinomial regression with PST Approach A and Avoid B as predictor variables.

Observed	Predicted			Correct (%)
	Smoker	Ex-smoker	Non-smoker	
Smoker	29	3	9	70.7
Ex-Smoker	17	5	7	17.2
Non-Smoker	13	4	18	51.4
Overall Correct (%)	56.2	11.4	32.4	49.5

5. 2. 3. Study 1 Discussion

Individual differences in reward learning predicted smoker status with moderate accuracy. Relative to non-smokers, smokers and ex-smokers had decreased learning from reward. Our results are concordant with those of Baker et al. (2011), in that our non-dependent (i.e., non-smoker) group showed higher reward learning in the PST compared with the dependent (i.e., smoker) group. In contrast, Potts, Bloom, Evans, and Drobles (2014), in a flanker task, reported that ex-smokers and smokers were more sensitive to rewards compared with non-smokers. Unlike Baker et al. (2011), our dependent group demonstrated increased learning from punishment relative to the non-dependent group. Butler et al. (2017), observed poor performance monitoring in smokers and found that post-punishment slowing correctly identified current smokers more so than former smokers (80% vs 60%).

Baker et al. (2011) examined PST performance in a sample considered ‘dependent’ due to their combined alcohol, tobacco and substance consumption. Our results were concordant with those of Baker et al., in that our non-dependent (i.e., never-smoker) group showed higher reward learning in the PST compared with the dependent (i.e., smoker) group. In contrast, Potts et al. (2014), in a flanker task, reported that ex-smokers and smokers were more sensitive to rewards

compared with non-smokers. Our results differ from Baker et al.'s (2011) research on poly-drug use, in that our dependent group demonstrated increased learning from punishment compared with the non-dependent group. In a study on alcohol dependent patients compared with healthy controls, Rustemeier et al. (2012) found no group differences with regard to reward and punishment learning on the PST. It is possible that differences between severity and type of substance use may influence reward and punishment learning.

Some of the contrast between the current findings and previous research may be attributable to phenotypic and methodological differences among studies. For example, Nestor et al. (2018b) included ex-smokers, abstinent for at least 12 months prior to testing. Potts et al. (2014) included only smokers who smoked over 10 cigarettes per day for the past year. The current study included smokers who were on average low in nicotine dependency, and abstinence was operationalized by the participant's self-reported smoking behaviour in the past 30 days. Carballo and López (2014) found increased length of abstinence in cocaine-dependent participants improved performance in response to negative feedback on a flanker task. Prolonged nicotine abstinence may similarly affect punishment sensitivity in the PST. Nestor et al. (2018b) used the Monetary Incentive Delay task, which focuses on gain and loss anticipation, while Potts et al. (2014) used a modified flanker task without feedback.

Many researchers (e.g., Baker, Piper, McCarthy, Majeskie, & Fiore, 2004; Blum et al., 2000; Koob, 2009) posit a negative affect addiction stage, involving avoidance of negative emotional after-effects of drug use. Lower levels of dopamine D2 receptor availability have been observed in chronic addiction. Lower dopamine levels have also been associated with increased learning from punishment (Frank et al., 2004). Martin, Cox, Brooks, and Savage (2014) showed that smokers were hyper-responsive to the anticipation of punishment. It is conceivable that our current smokers showed increased sensitivity to punishment due to decreased dopaminergic activity, and indeed

smoking heaviness (indexed by CO level) correlated with learning from punishment. This may also explain why punishment learning did not predict ex-smoker group status, as this group was no longer experiencing the negative affect stage of their former addiction.

In conclusion, these findings provide an insight into the effects of smoking status on reward and punishment learning using the PST. The results suggest that the PST has some utility in discriminating between smokers, ex-smokers, and non-smokers. These behavioural findings may be useful in understanding which smoking-cessation techniques are most effective, based on their use of positive and negative reinforcement.

5.3. Study 2

5.3.1. Methods

5.3.1.1. Participants

184 participants (97 female, 87 male), with a mean age of 28.38 (13.3), were recruited from Trinity College Dublin, and the University College Dublin. Eleven subjects with <55% accuracy on the AB pair in the training phase of the PST were removed from analysis, with a final sample of 75 non-smokers, 43 current smokers, 25 ex-smokers, and 30 vapers. 58 of these participants were included in the Chapter 2 study examining alcohol-use risk.

5.3.1.2. Materials

Smoking Status – Questionnaire

As in Study 1.

Smoking Status – Bioverification

Smoking group status was assessed objectively using exhaled carbon monoxide (CO) readings, and salivary cotinine (NicAlert strip) for a sub-set of participants. Current smokers were required to have a CO reading of at least 6 parts per million, and vapers, ex-smokers, and non-smokers were required

to have a CO reading less than 6ppm. Ex-smokers and non-smokers were required to have a NicAlert saliva reading of less than 2, while vapers were required to have a saliva reading of at least 1.

5.3.1.2.1. Machine Learning - Group Classification

A penalised multinomial logistic regression classifier implemented in sci-kit learn (Pedregosa et al., 2011) was used to predict smoking status group using the computational model of decision-making parameters, BIS, age, and gender. Specifically, mean model parameters from the training phase of the PST (a, v, t, alpha-pos, alpha-neg), test phase of the PST (a, t, vLL, vWL, vLL), BIS 2nd order subscales (attentional, motor, non-planning impulsivity), age, and gender were entered as model features. Features were first standardized/z-scored. A logistic regression model with five-fold cross-validation and elastic net penalty was fit to the data, the best-fitting model (hyper)parameters (I1 and C) from this model were then entered into a second logistic regression model and the data were re-fit using five-fold cross-validation. This process was repeated 100 times to improve precision/reduce variance of model estimates. Mean model accuracy metrics (MAE, MSE) and predictions from the re-fit logistic regression classifier were saved across all 100 iterations, and are presented below. Using the same pipeline, 100 null models with the same features and shuffled group labels were run as a comparison with the test models.

5.3.2. Results

5.3.2.1. Behavioral Results

Substance-use characteristics from the ESPAD questionnaire are presented in Table 5.3, and were available for 41/75 participants in the Non-smoker group, 35/43 participants in the Current Smoker group, and all participants in the Ex-smoker and Vaper groups. The Non-smoker group showed significantly lower alcohol intoxication, drug-use, and cannabis-use frequency compared with the Current Smoker group.

Table 5. 3. Median (IQR) substance-use characteristics for each sample group with between-groups significance test.

	Current Smokers (n = 35)	Ex-smokers (n = 25)	Vapers (n = 30)	Non-smokers (n= 41)	Kruskal Wallis (p)
Alcohol Intoxication Frequency					
<i>In the lifetime</i>	3 (6)	3.25 (6)	2 (6)	3 (6)	12.73 (.01)
<i>In the past 12 months</i>	3 (6)	2 (5)	2 (6)	2 (5)	12.25 (.01)
<i>In the past 30 days</i>	2 (3)	0 (2)	0 (3)	1 (3)	10.01 (.02)
Cannabis Frequency					
<i>In the lifetime</i>	4.5 (6)	5 (6)	5 (6)	2 (6)	23.19 (<.001)
<i>In the past 12 months</i>	3 (6)	1 (6)	2 (6)	1 (4)	15.68 (.001)
<i>In the past 30 days</i>	1 (5)	0 (5)	0 (6)	0 (4)	8.17 (.04)
Drugs Frequency (amphetamines, tranquillizers / sedatives, ecstasy, LSD, crack, cocaine, heroin, magic mushrooms" GHB, anabolic-steroids)					
<i>In the lifetime</i>	3.5 (6)	3 (6)	4 (6)	0 (6)	15.82 (.001)
<i>In the past 12 months</i>	2.5 (6)	1 (5)	0.75 (4)	0 (4)	11.54 (.01)

<i>In the past 30 days</i>	1 (4)	0 (2)	0 (1)	0 (1)	15.29 (.002)
----------------------------	-------	-------	-------	-------	--------------

Behavioral findings from the PST are presented below (Table 5. 3). In the test phase of the PST, Approach AC and BD did not differ significantly between groups ($p's > .12$). The Vaper group had lower choice accuracy for EF pairs in the PST Training Phase compared with Smokers and Non-smokers. In the Test phase, the Non-smoker group had higher mean choice selection of the highly rewarded stimuli (A, C) compared with current and ex-smokers, and the Vaper group had higher mean BD stimulus selections. However, these mean group differences did not survive multiple comparisons corrections (all $p's > .13$).

Table 5. 4. Mean (SD) accuracy in the PST training and test phase for each group with between-groups significance test.

	Current Smokers (n = 43)	Ex-smokers (n = 25)	Vapers (n = 30)	Non-smokers (n=75)	Kruskal Wallis (p)
PST Training Phase					
<i>A Accuracy</i>	81.16 (13.53)	81.9 (12.02)	80.83 (14.28)	83.17 (11.58)	0.62 (0.89)
<i>C Accuracy</i>	71.54 (18.64)	71.38 (13.97)	74.0 (16.04)	73.53 (16.21)	0.8 (0.85)
<i>E Accuracy</i>	58.18 (22.97)	52.05 (23.47)	47.0 (14.97)	54.58 (18.17)	7.64 (0.05)
PST Test Phase					
<i>Approach AC</i>	70.93 (16.56)	70.28 (17.53)	73.59 (17.28)	76.98 (15.88)	5.87 (0.12)
<i>Approach BD</i>	66.79 (20.52)	65.53 (19.37)	59.35 (14.38)	63.53 (18.86)	5.59 (0.13)

Trial RTs >4 seconds and <.15 seconds were removed from the group comparisons and are presented below. Mean RT during the Training phase different between groups, however corrected post-hoc group comparisons were non-significant ($p's > .09$). Group differences on CD pair trials

showed that ex-smokers had longer RTs compared with current smokers ($t = -2.96$, $p = 0.02$). On EF pair trials, vapers had longer RTs compared with current smokers ($t = -2.8$, $p = .03$). In the test phase, no significant group differences in RT were observed.

Table 5. 5. Mean (SD) reaction times in the PST training and test phase for each group with between-groups significance test.

	<i>Current Smokers</i>	<i>Ex-smokers</i>	<i>Vapers</i>	<i>Non-smokers</i>	<i>Kruskal Wallis (p value)</i>
<i>PSTWFB – All trials</i>	0.94 (0.34)	1.2 (0.42)	1.18 (0.49)	1.04 (0.42)	8.88 (0.03)
<i>AB</i>	0.9 (0.36)	1.1 (0.39)	1.09 (0.52)	0.95 (0.43)	6.49 (0.09)
<i>CD</i>	0.9 (0.35)	1.23 (0.47)	1.13 (0.48)	1.01 (0.45)	9.83 (0.02)
<i>EF</i>	1.02 (0.37)	1.29 (0.45)	1.32 (0.55)	1.16 (0.45)	9.22 (0.03)
<i>PSTNFB – All trials</i>	0.91 (0.35)	1.02 (0.38)	1.08 (0.52)	0.9 (0.35)	4.95 (0.18)
<i>AC Trials</i>	0.78 (0.47)	0.81 (0.68)	0.79 (0.41)	0.79 (0.45)	2.28 (0.51)
<i>BD Trials</i>	0.82 (0.89)	0.78 (0.27)	0.77 (0.22)	0.79 (0.47)	5.27 (0.15)
<i>LL</i>	1.12 (0.49)	1.29 (0.56)	1.29 (0.64)	1.1 (0.45)	3.79 (0.29)
<i>WL</i>	0.84 (0.31)	0.93 (0.37)	1.02 (0.5)	0.84 (0.34)	4.44 (0.22)
<i>WW</i>	0.86 (0.34)	0.98 (0.36)	1.03 (0.53)	0.85 (0.35)	5.35 (0.15)

A MANOVA comparing BIS sub-scales between smoking groups showed a significant effect ($F = 2.47$, $p = .009$) Post-hoc tests showed that the Vaper group had significantly lower attentional impulsivity compared with non-smokers ($t = 2.98$, $p(\text{bonf}) = .02$) and current smokers ($t = 3.09$, $p(\text{bonf}) = .01$). Current smokers had higher motor and non-planning impulsivity compared with ex-smokers and vapers, though these did not survive correction for multiple comparisons (p 's $> .26$).

Table 5. 6. Mean BIS scales scores (SD) for each group with between-groups significance test.

<i>BIS</i>	<i>Current Smokers</i>	<i>Ex-smokers</i>	<i>Vapers</i>	<i>Non-smokers</i>	<i>Kruskal Wallis (p value)</i>
<i>Attentional</i>	17.47 (3.97)	16.24 (3.09)	14.73 (3.51)	17.12 (3.83)	10.04 (0.02)
<i>Motor</i>	23.84 (4.37)	21.76 (3.52)	21.77 (4.05)	22.77 (4.49)	6.51 (0.09)
<i>Non-planning</i>	25.77 (5.24)	25.0 (5.57)	23.83 (5.45)	23.76 (4.75)	5.41 (0.14)

5. 3. 2. 2. Modelling of PST Training Phase

An RLHDDM model was run with 35,000 samples (5,000 discarded, with every 3rd sample retained).

The five parameters, a , t , v , pos_learn , and neg_learn were free to vary within the whole group.

Model convergence was assessed with the gelman-rubin r^{\wedge} statistic and visual inspection of the trace and autocorrelation plots (max. $r^{\wedge} = 1.04$, 99.66% of individual subject-level r^{\wedge} values were ≤ 1.02 , hence indicating adequate convergence). Model fit was assessed using the Deviance Information Criteria (DIC); 43908.52.

Posterior predictive checks were conducted by simulating task data from the posterior of the fitted model and comparing this to the observed data. A random sample from model trace was used to generate accuracy and RT data for each of the training phase stimulus pairs (AB, CD, EF), this step was repeated 100 times to better capture variability in the posterior distribution. These were then compared with the observed data. Figure 5.1 shows the observed versus simulated RT (averaged over 100 iterations) data for each stimulus types (AB, CD, EF). Error responses are negative on the x axis, and correct responses are positive. The simulated data predicted RTs that were very similar to those in the observed data. The simulated choice data showed some deviations from the observed data, with overpredicted mean choices in the EF pair.

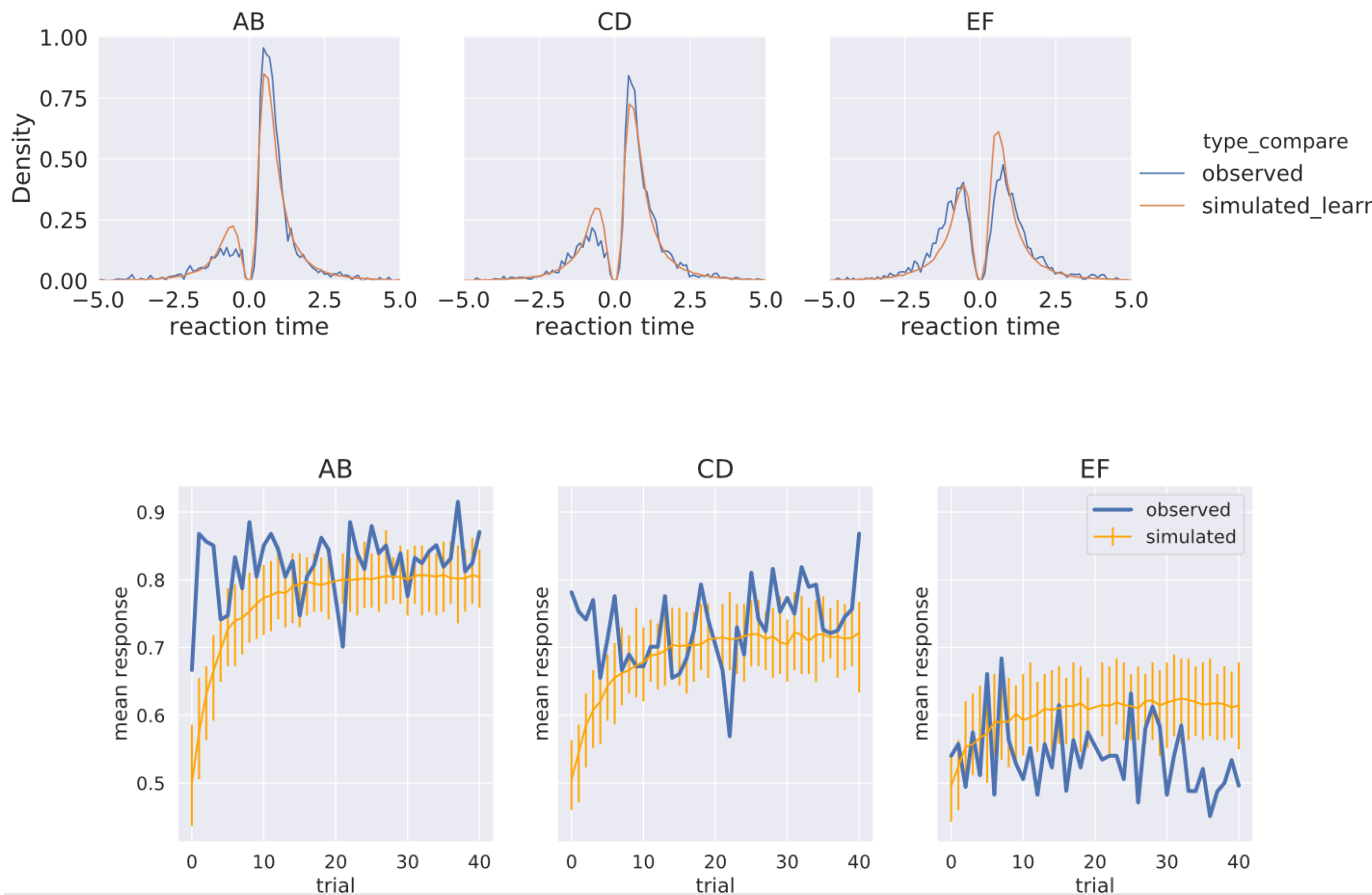


Figure 5. 1. Simulated and observed reaction times across stimulus types (upper), and simulated and observed mean choices (1=correct, 0=incorrect) across stimulus types (lower).

5. 3. 2. 3. Modelling of PST Test Phase

Two HDDM models were run with 12,000 samples and the first 2,000 discarded. Model 1 estimated non-decision time (t), boundary separation (a), and drift-rate (v) free to vary for the whole group. Model convergence was assessed with the gelman-rubin r^{\wedge} statistic and visual inspection of the trace and autocorrelation plots (max. $r^{\wedge} = 1.004$). Model fit was assessed using the Deviance Information Criteria (DIC); 39053.72. Model 2 estimated the same parameters, with drift-rate free to vary depending on the level of conflict in stimulus pairs (i.e., v_{LL} , v_{WL} , v_{WW}). LL stimulus pairs were those which included two shapes that were consistently punished in the training phase of the

PST (e.g., BD, DF), WW were those with two consistently rewarded shapes (e.g., AC, CE), and WL were those with a combination of rewarded and punished shapes (e.g., AD, CB). It was hypothesised that drift-rates for WL stimulus pairs would be lower, as the larger difference in value between shape stimuli would lead to a faster time to reach the correct choice boundary. Model 2 showed good convergence (max. rhat = 1.01), and better fit than Model 1 (DIC = 34089.42). The mean parameters from Model 2 are presented in Figure 5. 2. and Table 5. 6.

A MANOVA comparing computational model parameters between smoking status groups showed no significant differences between groups ($F(3,160) = 1.29, p = .14$). Kruskal Wallis tests comparing each parameter separately, showed a trend towards significance for the RLHDDM boundary separation parameter ($p = .03$) and the drift-rate for Win-Wins ($p = .03$). Post-hoc tests with bonferroni corrected p-values were conducted to explore these group differences. In the Training phase, the Vaper group showed significantly higher boundary separation values compared with current smokers ($t = -2.86, p = .029$). Ex-smokers showed larger non-decision times compared with non-smokers, however this did not survive the Bonferroni correction. In the Test phase, non-smokers showed higher drift-rates for win-win pairs compared with current and ex-smokers, and smaller non-decision times compared with vapers, however these comparisons were not significantly different ($p's > 0.17$).

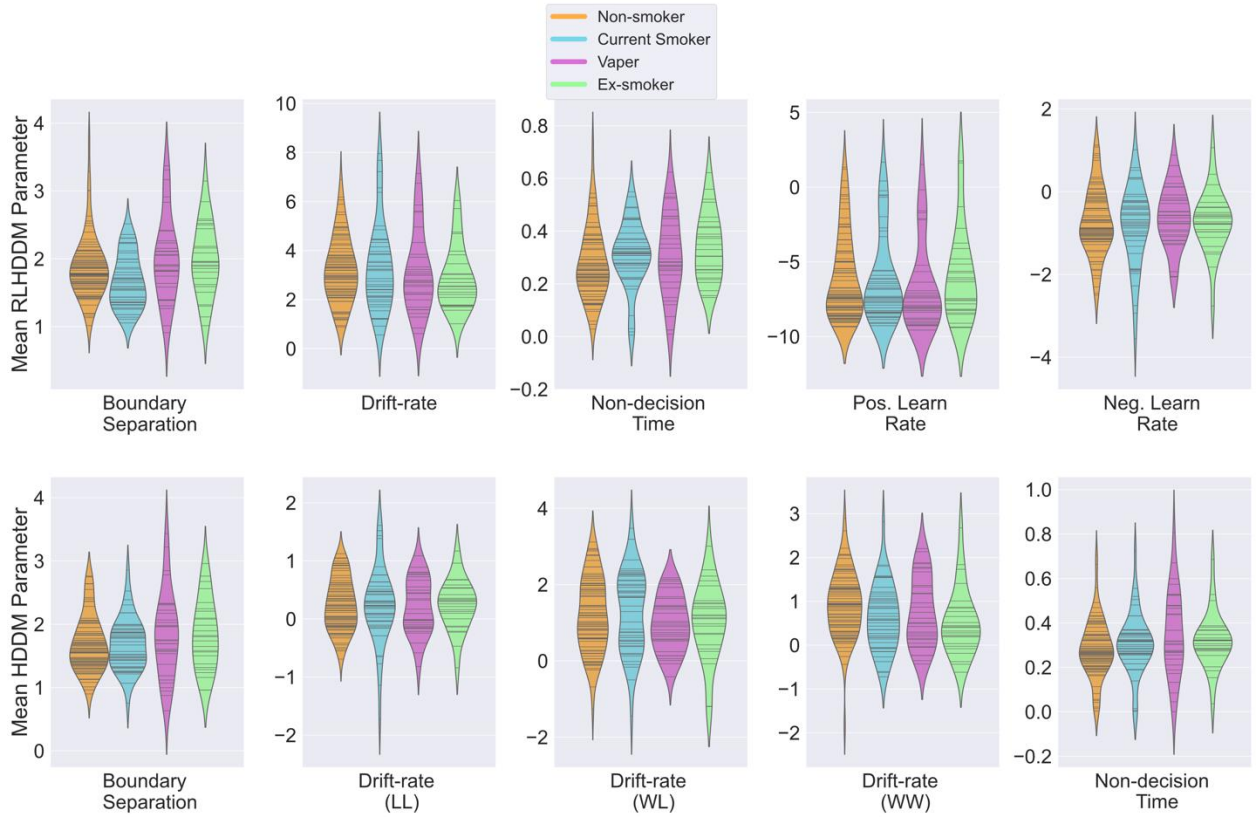


Figure 5. 2. Violin plots depicting mean RLHDDM and HDDM parameter estimates by smoking group.

Table 5. 7. Mean parameter estimates for each group, with significance test.

<i>Parameter</i>	Current Smokers	Ex-smokers	Vapers	Non-smokers	Kruskal Wallis (<i>p</i> value)
RLHDDM					
<i>Boundary Separation</i>	1.67 (0.39)	1.98 (0.53)	2.01 (0.64)	1.88 (0.47)	9.75 (0.02)
<i>Drift-rate</i>	3.34 (1.79)	2.91 (1.3)	3.16 (1.7)	3.23 (1.35)	1.61 (0.66)
<i>Non-decision Time</i>	0.31 (0.13)	0.34 (0.13)	0.31 (0.16)	0.28 (0.14)	5.89 (0.12)
<i>Neg. learning rate</i>	0.08 (0.18)	0.08 (0.23)	0.06 (0.17)	0.06 (0.17)	2.0 (.57)
<i>Pos. learning rate</i>	0.33 (0.17)	0.34 (0.15)	0.37 (0.15)	0.34 (0.17)	0.79 (0.85)
HDDM					
<i>Boundary Separation</i>	1.68 (0.42)	1.8 (0.56)	1.83 (0.56)	1.7 (0.45)	2.02 (0.56)
<i>Drift-rate (Loss-Loss)</i>	0.23 (0.65)	0.241 (0.44)	0.21 (0.48)	0.24 (0.45)	0.15 (0.98)
<i>Drift-rate (Win-Loss)</i>	1.27 (1.06)	1.02 (1.0)	0.97 (0.73)	1.19 (0.97)	2.03 (0.57)
<i>Drift-rate (Win-Win)</i>	0.62 (0.76)	0.52 (0.76)	0.76 (0.8)	0.92 (0.76)	8.59 (0.035)
<i>Non-decision Time</i>	0.31 (0.14)	0.32 (0.13)	0.35 (0.19)	0.28 (0.13)	4.44 (0.23)

5. 3. 2. 4. Group Classification: Current Smokers vs. Non-smokers

The PST Summary Model mean Brier score (0.22, SD = .006) was significantly lower than the baseline model (0.23, .005); $t(198) = -7.69, p < .001$. The model AUC (0.61, .03) outperformed the baseline model (0.56, .04) in 81% of iterations; $t(198) = 11.01, p < .001$, and the F1 score (0.28, .009) outperformed the baseline (0.27, .008) in 79% of iterations; $t(198) = 10.21, p < .001$. All features in the model were significant predictors of smoking group.

Table 5. 8. PST Summary Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Demographic</i>	<i>Mean Beta - Demographic</i>
<i>intercept</i>	5	-0.585	5	0.56
<i>Age</i>	4.93	0.308	4.35	0.21
<i>Approach AC</i>	4.86	-0.305	2.11	-0.002
<i>Approach BD</i>	4.12	0.142	2.23	0.006
<i>Gender</i>	4.05	-0.129	2.64	-0.06

The Computational Model mean Brier score (0.227, .009) was significantly lower than the baseline model (0.235, .006); $t(198) = -7.28, p < .001$. The model AUC (0.61, .04) outperformed the baseline model (0.52, .05) in 82% of iterations; $t(198) = 12.9, p < .001$, and the F1 score for the test model (0.28, .008) outperformed the baseline model (0.27, .006) in 78% of iterations; $t(198) = 9.26, p < .001$. Age, gender, boundary separation from the RLHDDM, drift-rate for win-loss and win-win conditions, non-decision time from the HDDM, and positive learning rate were significant predictors.

Table 5. 9. Computational Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Null</i>	<i>Mean Beta - Null</i>
<i>intercept</i>	5	-0.596	5	-0.57
<i>RLHDDM: Boundary Separation</i>	4.54	-0.332	1.14	-0.002
<i>Age</i>	4.47	0.264	3.58	0.14
<i>HDDM: Drift-rate Win-Win</i>	4.28	-0.197	1.29	0.005
<i>RLHDDM: Non-decision Time</i>	3.78	0.186	1.29	0.002
<i>Gender</i>	2.43	-0.067	1.56	-0.03
<i>HDDM: Non-decision Time</i>	1.68	0.011	1.29	-0.003
<i>HDDM: Drift-rate Win-Lose</i>	1.56	0.039	1.37	-0.009
<i>RLHDDM: Pos. Learning Rate</i>	1.48	-0.015	1.34	0.003
<i>HDDM: Drift-rate Lose-Lose</i>	1.19	-0.009	1.31	0.002
<i>RLHDDM: Neg. Learning Rate</i>	1.07	-0.009	1.35	-0.006
<i>HDDM: Boundary Separation</i>	1.06	0.033	1.33	0.004
<i>RLHDDM: Drift Rate</i>	0.99	-0.019	1.29	-0.009

The Personality Model mean Brier score (0.229, SD = .004) was significantly lower than the baseline model (0.231, .005); $t(198) = -4.62, p < .001$. The model AUC (0.58, .03) outperformed the baseline model (0.55, .04) in 71% of iterations; $t(198) = 6.51, p < .001$, and the model F1 score (0.279, .009) outperformed the baseline (0.273, .007) in 68% of iterations; $t(198) = 5.16, p < .001$. Age, gender, and non-planning impulsivity were significant features.

Table 5. 10. Personality Model mean choice frequencies and beta values for each feature in the test and demographic model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Null</i>	<i>Mean Beta - Null</i>
<i>intercept</i>	5	-0.569	5	-0.57
<i>Age</i>	4.47	0.196	4.2	0.21
<i>BIS Non-planning</i>	4.18	0.151	1.95	0.001
<i>Gender</i>	2.62	-0.049	2.46	-0.05
<i>BIS Motor</i>	1.84	0.021	1.98	0.003
<i>BIS Attentional</i>	0.81	-0.007	2.05	0.005

The Combined Model Brier score (0.229, SD = .008) was also significantly lower than the baseline model (0.232, .006); $t(198) = -2.89, p = .004$. The mean model AUC (0.59, .04) outperformed the baseline model (0.53, .05) in 73% of iterations; $t(198) = 8.65, p < .001$, and the model F1 score (0.279, .009) outperformed the baseline (0.273, .009) in 72% of iterations; $t(198) = 4.86, p < .001$. In the Combined Model, age, gender, non-planning impulsivity, boundary separation (RLHDDM), drift-rate for win-win conditions, non-decision time (RLHDDM & HDDM), and positive learning rate were significant features.

Table 5. 11. Combined Model mean choice frequencies and beta values for each feature in the test and demographic model.

Feature	Mean choice frequency - Test	Mean Beta - Test	Mean choice frequency - Null	Mean Beta - Null
<i>intercept</i>	5	-0.595	5	-0.57
<i>Age</i>	4.31	0.219	3.28	0.12
<i>RLHDDM: Boundary Separation</i>	4.3	-0.286	1	0.001
<i>BIS: Non-planning</i>	3.96	0.148	1.1	-0.002
<i>HDDM: Drift-rate Win-Win</i>	3.9	-0.159	1.03	-0.003
<i>RLHDDM: Non-decision Time</i>	3.28	0.138	0.85	0.003
<i>Gender</i>	1.99	-0.057	1.13	-0.018
<i>HDDM: Non-decision Time</i>	1.26	0.007	1.02	-0.0009
<i>RLHDDM: Pos. Learning Rate</i>	1.22	-0.007	1.11	-0.0004
<i>BIS: Motor</i>	1.1	0.011	0.85	0.006
<i>HDDM: Drift-rate Win-Lose</i>	1.02	0.024	1.03	-0.005
<i>RLHDDM: Neg. Learning Rate</i>	0.9	-0.016	1.07	-0.01
<i>HDDM: Boundary Separation</i>	0.86	0.033	1.1	-0.002
<i>HDDM: Drift-rate Lose-Lose</i>	0.8	0.005	1.21	-0.007
<i>RLHDDM: Drift-rate</i>	0.73	-0.017	1.04	-0.003
<i>BIS: Attentional</i>	0.61	-0.009	0.87	-0.008

Machine learning metrics (Brier & AUC scores) for the binary logistic regression with smoking group (non-smoker vs. current smoker) as the dependent variable are presented below. The four models showed similar fit.



Figure 5. 3. Violin plot of machine learning scores for binary classification.

5. 3. 2. 5. Multinomial Group Classification – Current Smokers, Vapers, Ex-smokers, Non-smokers

The PST Model mean balanced accuracy score (0.36, .02) was significantly higher than the baseline model (0.27, .03); $t(198) = -29.88, p < .001$. The test model Cohen’s Kappa score (0.18, .04) was also significantly higher than the baseline model score (.03, .04); $t(198) = -29.73, p < .001$, and the Matthew’s correlation coefficient (test model = 0.18 (.04), baseline model = .04 (.04); $t(198) = -30.0, p < .001$). All features in the test model were significant, with mean feature selections $> 95^{\text{th}}$ percentile of demographic model feature selections.

Table 5. 12. Mean selection of features across five folds, demographic model selections are in parentheses.

Feature	Non-smoker	Current Smoker	Ex-smoker	Vaper
<i>Approach AC</i>	4.34 (1.05)	4.25 (1.08)	4.33 (1.16)	4.12 (1.07)
<i>Approach BD</i>	4.34 (1.08)	4.29 (1.08)	4.4 (1.17)	4.08 (1.12)
<i>Age</i>	4.36 (1.08)	4.32 (1.08)	4.4 (1.2)	4.16 (1.12)
<i>Gender</i>	4.3 (1.06)	4.26 (1.05)	4.32 (1.19)	4.1 (1.11)

Table 5. 13. Mean beta values for each feature by smoking group for the test model, demographic model betas are in parentheses.

Feature	Non-smoker	Current Smoker	Ex-smoker	Vaper
<i>Approach AC</i>	0.433 (0.046)	-0.392 (-0.081)	-0.432 (0.226)	0.078 (-0.145)
<i>Approach BD</i>	-0.214 (0.099)	0.239 (0.245)	0.308 (-0.187)	-0.267 (-0.457)
<i>Age</i>	-0.497 (-0.553)	-0.831 (-0.825)	0.439 (0.346)	0.775 (0.962)
<i>Gender</i>	-0.419 (-0.458)	0.326 (0.387)	-0.073 (0.015)	0.431 (0.428)

Table 5. 14. Classification table for the PST Summary model, with percentage of correct group classifications.

<i>Observed</i>	<i>Predicted</i>				Correct (%)
	Non-smoker	Current Smoker	Ex-smoker	Vaper	
Non-smoker	37.16	22.41	6.93	8.5	49.55
Current Smoker	13.63	19.05	4.99	5.33	44.30
Ex-smoker	5.37	10.11	2.33	7.19	9.32
Vaper	6.37	7.11	3.72	12.8	42.67
Overall (%)	36.14	33.91	10.39	19.55	41.24

The Computational parameter model mean balanced accuracy score (0.36, SD = .03) outperformed the demographic model (0.32, SD = .02); $t(198) = -11.35$, $p < .001$. This was also the case for the Cohen's Kappa score (test model = 0.18 (.03), baseline model = 0.11 (.03)); $t(198) = -13.02$, $p < .001$., and for Matthews correlation coefficient (test model = 0.18 (0.3), baseline model = 0.11 (.03)); $t(198) = -13.09$, $p < .001$. Age was a significant feature for all groups, gender was significant for all groups except the ex-smoker group, non-decision time (RLHDDM) was significant features for the non-smoker, current smoker, and ex-smoker groups, boundary separation (RLHDDM) was significant for the non-smoker and smoker groups. Drift-rate for the Loss-Loss condition (HDDM) was significant for the non-smoker group, and drift-rate for the Win-Win condition (HDDM) and non-decision time (HDDM) were significant for the Vaper group.

Table 5. 15. Mean feature selection across test model folds for each group in the Computational test model with Demographic model selections in parentheses.

<i>Feature</i>	<i>Non-smoker</i>	<i>Current Smoker</i>	<i>Ex-smoker</i>	<i>Vaper</i>
HDDM				
<i>Boundary Separation</i>	4.26 (4.15)	4.44 (3.94)	4.21 (4.11)	4.32 (4.27)
<i>Drift-rate (Loss-Loss)</i>	4.45 (4.38)	4.24 (4.39)	4.69 (4.17)	4.47 (4.12)
<i>Drift-rate (Win-Loss)</i>	4.2 (4.03)	4.37 (3.58)	4.3 (3.85)	4.63 (3.96)
<i>Drift-rate (Win-Win)</i>	4.78 (4.28)	4.75 (4.39)	4.75 (4.11)	4.23 (4.0)
<i>Non-decision Time</i>	4.14 (3.79)	4.07 (4.13)	4.56 (4.14)	4.72 (3.56)
RLHDDM				
<i>Boundary Separation</i>	4.78 (4.39)	4.78 (4.29)	4.22 (4.35)	4.35 (4.11)
<i>Drift-rate</i>	4.15 (4.33)	4.37 (4.02)	4.35 (4.22)	4.44 (3.9)
<i>Non-decision Time</i>	4.77 (3.61)	4.78 (3.75)	4.78 (3.74)	4.19 (4.08)
<i>Pos. Learning Rate</i>	4.09 (4.13)	4.24 (4.32)	4.18 (4.34)	4.48 (4.02)
<i>Neg. Learning Rate</i>	4.09 (4.06)	4.36 (4.07)	4.31 (3.99)	4.25 (3.83)
Demographic				
<i>Age</i>	4.78 (4.39)	4.78 (4.39)	4.78 (4.37)	4.78 (4.39)
<i>Gender</i>	4.78 (4.39)	4.78 (4.39)	4.57 (4.01)	4.78 (4.39)

Table 5. 16 Mean betas for each feature and group in the Computational test model, with demographic model betas in parentheses.

<i>Feature</i>	<i>Non-smoker</i>	<i>Current Smoker</i>	<i>Ex-smoker</i>	<i>Vaper</i>
HDDM				
<i>Boundary Separation</i>	-0.165 (-0.306)	0.253 (0.277)	0.037 (-0.228)	-0.124 (0.335)
<i>Drift-rate (Loss-Loss)</i>	-0.088 (0.249)	-0.071 (-0.333)	0.241 (-0.179)	0.151 (0.194)
<i>Drift-rate (Win-Loss)</i>	-0.137 (-0.144)	0.139 (0.043)	-0.049 (0.133)	-0.302 (-0.128)
<i>Drift-rate (Win-Win)</i>	0.431 (-0.207)	-0.229 (0.366)	-0.527 (-0.186)	0.106 (0.044)
<i>Non-decision Time</i>	-4.79E-06 (0.085)	-0.070 (-0.159)	-0.276 (0.218)	0.296 (-0.059)
RLHDDM				
<i>Boundary Separation</i>	0.431 (0.535)	-0.912 (-0.304)	0.0361 (-0.378)	0.204 (-0.280)
<i>Drift-rate</i>	0.086 (0.259)	-0.075 (-0.178)	-0.063 (-0.234)	0.124 (-0.065)
<i>Non-decision Time</i>	-0.360 (-0.011)	0.382 (0.001)	0.482 (-0.039)	-0.121 (-0.108)
<i>Pos. Learning Rate</i>	0.059 (-0.128)	0.031 (0.163)	0.092 (-0.260)	-0.147 (0.102)
<i>Neg. Learning Rate</i>	0.046 (-0.148)	-0.033 (0.119)	0.065 (0.127)	-0.095 (0.146)
Demographic				
<i>Age</i>	-0.536 (-0.414)	-0.737 (-0.831)	0.491 (0.406)	0.756 (0.848)
<i>Gender</i>	-0.463 (-0.379)	0.397 (0.333)	0.002 (-0.115)	0.387 (0.395)

Table 5. 17. Classification table of Computational Model results, with mean classification and percentages across 100 model iterations.

<i>Observed</i>	<i>Predicted</i>				Correct (%)
	Non-smoker	Current Smoker	Ex-smoker	Vaper	
Non-smoker	39.28	19.94	9.63	6.15	52.37%
Current Smoker	12.96	20.37	5.4	4.27	47.37%
Ex-smoker	6.76	8.53	3.07	6.64	12.28%
Vaper	6.75	5.25	8.49	9.51	31.7%

The mean Personality Model balanced accuracy score was higher for the test model (0.41, .02) compared with the null model (0.31, .02); $t(198) = -25.39, p < .0001$. The Matthew's correlation coefficient score was significantly higher for the test model (0.23, .03) compared with the null model (.08, .03); $t(198) = -26.26, p < .0001$, as was the Cohen Kappa score (test model = 0.23 (.03), null model = 0.08 (.03); $t = -26.35, p < .0001$). All features in the test model were chosen > 95th of selections in the demographic model.

Table 5. 18. Mean feature selections for the Personality Model, with demographic model selections in parentheses.

<i>Feature</i>	<i>Non-smoker</i>	<i>Current Smoker</i>	<i>Ex-smoker</i>	<i>Vaper</i>
<i>BIS: Attentional</i>	4.84 (2.6)	4.55 (2.59)	4.6 (2.61)	4.84 (2.76)
<i>BIS: Motor</i>	4.47 (2.79)	4.64 (2.67)	4.83 (2.7)	4.68 (2.83)
<i>BIS: Non-planning</i>	4.84 (2.71)	4.82 (2.61)	4.8 (2.6)	4.7 (2.83)
<i>Age</i>	4.84 (2.83)	4.84 (2.83)	4.84 (2.83)	4.84 (2.83)
<i>Gender</i>	4.84 (2.83)	4.84 (2.83)	4.7 (2.72)	4.84 (2.83)

Table 5. 19. Mean beta values for each feature in the Personality Model, with demographic model betas in parentheses.

<i>Feature</i>	<i>Non-smoker</i>	<i>Current Smoker</i>	<i>Ex-smoker</i>	<i>Vaper</i>
<i>BIS: Attentional</i>	0.353 (0.049)	0.066 (-0.070)	-0.144 (-0.048)	-0.669 (0.223)
<i>BIS: Motor</i>	0.120 (-0.261)	0.114 (0.088)	-0.457 (0.021)	-0.181 (0.363)
<i>BIS: Non-planning</i>	-0.495 (0.140)	0.238 (0.056)	0.387 (0.032)	0.261 (-0.447)
<i>Age</i>	-0.409 (-0.481)	-0.815 (-0.836)	0.269 (0.350)	0.782 (0.862)
<i>Gender</i>	-0.421 (-0.415)	0.362 (0.356)	0.032 (-0.022)	0.419 (0.371)

Table 5. 20. Classification table for the Personality model, with percentage of correct group classifications.

<i>Observed</i>	<i>Predicted</i>				Correct (%)
	Non-smoker	Current Smoker	Ex-smoker	Vaper	
Non-smoker	38.8	20.41	8.57	7.22	51.73
Current Smoker	15.31	19.51	3.44	4.74	45.37
Ex-smoker	6.99	7.03	2.97	8.01	11.88
Vaper	2.45	6.46	4.37	16.72	55.73
Overall (%)	37.73	30.87	11.19	21.21	45.09

The Combined Model mean balanced accuracy score was 0.379 (.02), and did not differ significantly from the demographic model (0.376, SD = .03); $t(198) = -0.769$, $p = 0.44$. The overall accuracy score, not adjusted for group sizes, was significantly higher for the test model (0.43, SD = .02) than the demographic model (0.41, SD = .03); $t(198) = -4.61$, $p < .0001$. The Cohen's Kappa score was significantly higher for the test model (0.19, SD = .03) compared with the demographic model (0.18, SD = .04); $t(198) = 4.25$, $p < .0001$. The Matthew's correlation coefficient was also higher for the test model (0.21, SD = .03) compared with the demographic model (0.18, SD = .04); $t(198) = -4.182$, $p < .0001$.

Figure 5. 21. Mean selections of features across five cross-validation folds, demographic model selections are in parentheses.

<i>Feature</i>	<i>Non-smoker</i>	<i>Current Smoker</i>	<i>Ex-smoker</i>	<i>Vaper</i>
<i>HDDM</i>				
<i>Boundary Separation</i>	4.77 (2.73)	4.64 (4.81)	4.29 (3.33)	4.39 (4.9)
<i>Drift-rate (Loss-Loss)</i>	4.62 (3.14)	4.42 (4.9)	4.96 (4.88)	4.76 (3.3)
<i>Drift-rate (Win-Loss)</i>	4.35 (3.13)	4.26 (4.73)	4.45 (3.35)	4.88 (4.9)
<i>Drift-rate (Win-Win)</i>	4.97 (3.24)	4.86 (4.9)	4.97 (4.7)	4.36 (3.82)
<i>Non-decision Time</i>	4.23 (4.11)	4.21 (3.1)	4.57 (4.0)	4.91 (4.08)
<i>RLHDDM</i>				
<i>Boundary Separation</i>	4.97 (4.39)	4.97 (4.7)	4.33 (4.79)	4.12 (3.3)
<i>Drift-rate</i>	4.38 (4.24)	4.42 (4.9)	4.52 (4.71)	4.45 (2.97)
<i>Non-decision Time</i>	4.94 (4.78)	4.96 (4.9)	4.97 (4.9)	4.42 (4.48)
<i>Pos. Learning Rate</i>	4.32 (3.94)	4.38 (3.3)	4.63 (3.89)	4.15 (3.58)
<i>Neg. Learning Rate</i>	4.35 (4.9)	4.62 (4.82)	4.48 (4.77)	4.39 (4.8)
<i>Personality</i>				
<i>BIS: Attentional</i>	4.97 (3.87)	4.42 (3.1)	4.71 (3.79)	4.97 (3.91)
<i>BIS: Motor</i>	4.55 (3.87)	4.55 (4.58)	4.95 (4.8)	4.41 (4.9)
<i>BIS: Non-planning</i>	4.97 (4.76)	4.86 (3.94)	4.93 (4.75)	4.85 (3.65)
<i>Demographic</i>				
<i>Age</i>	4.97 (4.9)	4.97 (4.9)	4.97 (4.88)	4.97 (4.9)
<i>Gender</i>	4.97 (4.9)	4.97 (4.9)	4.65 (3.98)	4.97 (4.9)


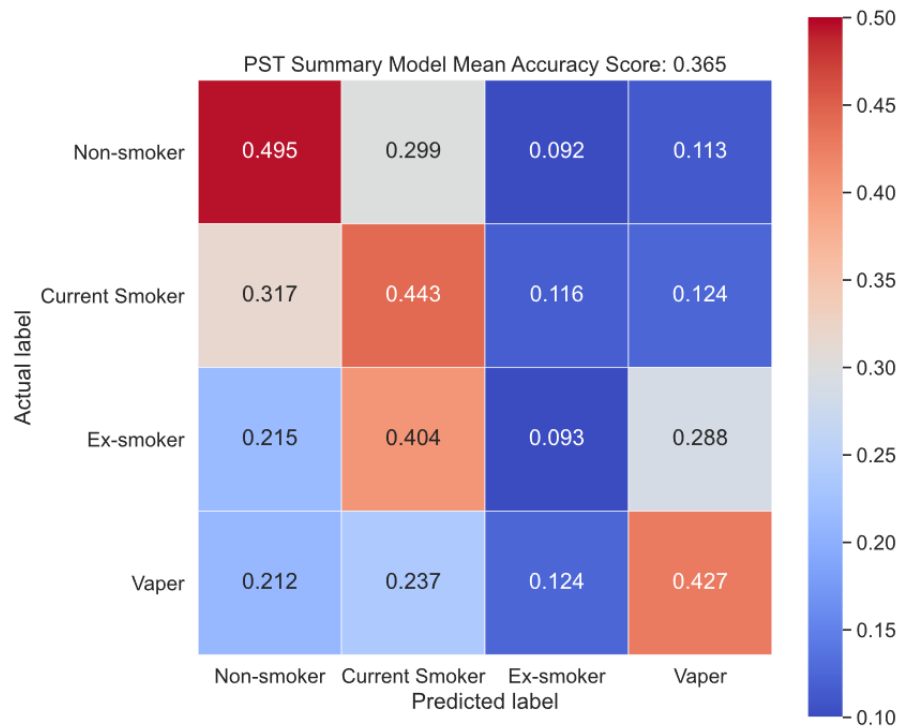
Figure 5. 22. Mean beta for each feature, demographic model selections are in parentheses.



<i>Feature</i>	<i>Non-smoker</i>	<i>Current Smoker</i>	<i>Ex-smoker</i>	<i>Vaper</i>
HDDM				
<i>Boundary Separation</i>	-0.244 (-0.060)	0.298 (-0.270)	0.009 (-0.003)	-0.069 (0.478)
<i>Drift-rate (Loss-Loss)</i>	-0.130 (-0.050)	-0.029 (0.317)	0.315 (-0.323)	0.206 (-0.055)
<i>Drift-rate (Win-Loss)</i>	-0.112 (0.034)	0.091 (0.240)	-0.067 (0.080)	-0.347 (-0.402)
<i>Drift-rate (Win-Win)</i>	0.502 (-0.029)	-0.204 (0.263)	-0.568 (-0.277)	-0.051 (0.033)
<i>Non-decision Time</i>	-0.012 (-0.092)	-0.059 (-0.038)	-0.248 (0.161)	0.314 (-0.155)
RLHDDM				
<i>Boundary Separation</i>	0.452 (0.177)	-0.914 (-0.245)	0.066 (0.252)	0.096 (-0.202)
<i>Drift-rate</i>	0.125 (0.151)	-0.098 (-0.388)	-0.061 (0.232)	0.054 (0.011)
<i>Non-decision Time</i>	-0.358 (-0.212)	0.355 (0.437)	0.477 (-0.402)	-0.139 (0.194)
<i>Pos. Learning Rate</i>	0.074 (0.111)	-0.057 (-0.060)	0.116 (-0.129)	0.005 (0.054)
<i>Neg. Learning Rate</i>	0.065 (-0.383)	-0.021 (0.226)	0.052 (0.239)	-0.071 (0.248)
Personality				
<i>BIS: Attentional</i>	0.493 (0.093)	0.083 (0.035)	-0.164 (-0.109)	-0.776 (-0.106)
<i>BIS: Motor</i>	0.126 (0.095)	0.106 (-0.176)	-0.490 (-0.339)	-0.141 (0.429)
<i>BIS: Non-planning</i>	-0.573 (-0.190)	0.232 (-0.096)	0.467 (0.325)	0.302 (-0.019)
Demographic				
<i>Age</i>	-0.465 (-0.367)	-0.743 (-0.745)	0.426 (0.381)	0.757 (0.690)
<i>Gender</i>	-0.478 (-0.388)	0.425 (0.290)	0.065 (-0.044)	0.391 (0.434)



Table 5. 23. Classification table of Combined Model results, with mean classification and percentages across 100 model iterations.

<i>Observed</i>	<i>Predicted</i>				Correct (%)
	Non-smoker	Current Smoker	Ex-smoker	Vaper	
<i>Non-smoker</i>	39.4	18.31	9.25	8.04	52.53
<i>Current Smoker</i>	13.68	18.7	5.7	4.92	43.49
<i>Ex-smoker</i>	4.69	7.68	4.38	8.25	17.52
<i>Vaper</i>	5.16	4.99	8.34	11.51	38.37
<i>Overall (%)</i>	37.38	28.72	15.99	18.91	42.77

Correlation matrices depicting percentage of correct group classifications for the PST Summary, Computational, Personality, and Combined Models with corresponding demographic models are presented in Figure 5. 4.










Figure 5. 4. Confusion matrix heatmaps for each test and demographic model, with classification accuracy expressed as the percentage of true group classifications.

5. 3. 2. 6. Correlations between ML Features

Correlations between self-reported impulsivity, PST test phase performance, and computational parameter estimates showed consistent relationships between training and test phase DDMs, and relationships between self-report and task measures (highlighted in Table 5. 23.). Boundary separation and non-decision time parameters from the RLHDDM and HDDM showed a significant positive relationship, and a significant negative relationship was found between positive and negative learning rates. Significant correlations were also observed between drift-rates in the HDDM and behavioural test phase PST performance. Motor impulsivity showed a significant negative relationship with boundary separation from the RLHDDM, and a positive relationship with drift-rate from the RLHDDM. Drift-rate for Win-Losses in the HDDM showed a positive correlation with non-planning impulsivity.

Feature	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1. BIS Attention	—													
2. BIS Motor	0.423***	—												
3. BIS Non-planning	0.452***	0.606***	—											
4. RLHDDM: a	0.028	-0.209*	-0.164	—										
5. RLHDDM: v	-0.089	0.185*	0.179	-0.256**	—									
6. RLHDDM: t	0.092	0.013	0.033	0.286**	0.039	—								
7. RLHDDM: Neg. Learning	0.148	-0.036	-0.019	0.309**	-0.409***	0.088	—							
8. RLHDDM Pos. Learning	-0.083	-0.074	-0.123	0.056	0.034	-0.059	-0.265**	—						
9. HDDM a	0.105	-0.163	-0.169	0.739***	-0.169	0.309**	0.294**	0.021	—					
10. HDDM vWW	-0.160	0.036	0.118	0.003	0.282**	-0.101	-0.187	0.254**	-0.182	—				
11. HDDM vLL	-0.012	-0.026	-0.047	-0.066	0.097	-0.089	-0.053	0.079	-0.072	0.115	—			
12. HDDM vWL	-0.041	0.144	0.199*	-0.001	0.515***	0.072	-0.087	0.012	-0.019	0.449***	0.375***	—		
13. HDDM: t	0.031	-0.057	-0.039	0.215	-0.139	0.610***	0.161	-0.124	0.223*	-0.221*	-0.091	-0.059	—	
14. Approach AC	-0.080	0.018	0.118	0.233*	0.322**	0.093	-0.052	0.151	0.139	0.798***	0.150	0.718***	-0.058	—
15. Approach BD	0.061	0.037	0.035	0.125	-0.269**	0.112	0.066	0.023	0.182	0.048	0.709***	0.692***	-0.040	0.316***

*p < .05, ** p < .01, *** p < .0001

Table 5. 24. Partial Spearman's correlations between machine learning features, with age as a covariate. Significant correlations are highlighted between (i) self-report measures (orange), (ii) task-related parameters and variables (blue), (iii) task-related and self-report measures (green).

5.3.3. Study 2 - Discussion

The findings of Study 1 were extended by applying DDMs to the PST and comparing the relative fit of four penalized regression models with varying behavioural, computational, and personality features to predict smoking group status in a larger sample. Given the rise in electronic cigarette/vaping use (Hammond et al., 2019), a group of Vapers was included in the current analysis.

The binary models predicting smoker vs. non-smoker status outperformed demographic models. In the Combined Model predicting smoker versus non-smoker group status, increases in age, non-planning impulsivity, non-decision times, and reductions in boundary separation, positive learning rate, and drift-rate for Win-Wins significantly predicted smoking group membership. In the PST Summary model, reduced learning from positive and negative feedback predicted smoker group membership.

The multinomial models predicted group membership beyond chance levels (25%), and all models outperformed demographic models with the exception of the balanced accuracy score of the Combined Model. Ex-smokers were the most typically misclassified. Significant features varied according to group and conformed overall to the previous binary analysis for smoker and non-smoker groups. Age and non-decision time were the most frequently selected in the combined and computational multinomial models. The ex-smoker group was predicted by increased drift-rates for Loss-Loss condition in the PST test phase, and decreased drift-rate in the Win-Win condition. Increased non-decision times and motor impulsivity, and decreased non-planning impulsivity also predicted ex-smoker group status. Vaper group status was predicted by increased age and non-decision times, as well as decreased attentional impulsivity. Gender was also a significant predictor, with males increasing the odds of current and vaper group status, and reducing the odds of non-smoker group status.

The finding that smaller drift-rates predicted smoking group status supports the hypothesis that reductions in evidence accumulation are a transdiagnostic risk-factor relevant to substance-use (Sripada & Weigard, 2021; Weigard & Sripada, 2021). Weigard et al. (2021) found that reductions in drift-rate during a Go/No-Go task predicted prospective substance dependence, which included frequency of cigarette consumption. The results presented in the current study also show a relationship between reduced drift-rates and smoking status. Reduced drift-rate for Win-Wins was a significant predictor of ex-smoker group status, however *increases* in drift-rate for the Loss-Loss condition was also a significant predictor. It is not immediately apparent why divergent drift-rates were observed for the two high conflict conditions in the PST test phase. One possibility is that the contingencies of the most-rewarded stimuli were learned more accurately than the most-punished stimuli in the PST training phase, leading to faster accumulation of evidence for pairs containing these stimuli in the test phase. As drift-rate for Win-Wins predicted both smoker and ex-smoker group status, this suggests that impairments in evidence accumulation may persist following nicotine abstinence.

The PST Summary models suggested a pattern of results similar to Baker et al. (2011; 2013), who found that substance-dependent individuals showed less accuracy overall on the PST test phase. This was reflected in our group classifications; current smoker status was predicted by lower approach AC, and higher approach BD- indicating low utilization of the stimulus reward/punishment contingencies learned during the training phase. The ex-smoker group were also classified by low approach AC/BD accuracy, whereas the non-smoker and Vaper groups were classified by the reverse pattern.

Higher self-reported non-planning impulsivity, unlike attentional and motor impulsivity, predicted smoker group status relative to non-smoker. A number of studies have investigated multidimensional aspects of impulsivity in smokers, however there is no strong consensus on what aspects of trait impulsivity are most predictive of smoking status. Chase and Hogarth

(2011) found low to moderate relationships between impulsivity (as measured by the BIS) and smoking dependency and consumption. Of the three sub-scales, non-planning showed the strongest correlation with cigarette dependency. Ryan et al. (2013) found that attentional and non-planning impulsivity correlated with dependency, but not consumption, and Round et al. (2010) reported correlations between motor and non-planning impulsivity and cigarette dependence and frequency of consumption.

A limitation of the current study is its cross-sectional design, which does not allow for the investigation of changes associated with nicotine abstinence across time. Further, the length of abstinence among ex-smokers was not controlled for in the current study. A general limitation of the current analysis is that co-occurrence of non-nicotine substance-use was not controlled for. Young adult e-cigarette users have been found to engage in greater co-substance use compared with binge drinkers or cannabis users (Lanza et al., 2020), which may influence the extent of decision-making impairments observed. While the current study examined various smoking states, it did not control for alcohol or poly-drug use. Heavy drinking and smoking together produced an additive effect on delayed reward discounting in comparison with heavy drinking-only or smoking-only groups in a study by Moallem and Ray (2012), and it is possible that high-risk alcohol users in the current sample would show differential performance on the PST compared with non-risk users. Future research may seek to isolate various groups of high versus low risk users across a range of substances, or indeed apply clustering methods to identify sub-types of task-related behavioral and computational parameters related to substance-use.

5. 4. General Conclusion

This study demonstrates the utility of task-based summary and computational indices of value-based decision-making as predictors of nicotine-related substance-use. Study 1 demonstrated that the PST classified current smokers, ex-smokers, and non-smokers with moderate accuracy, with non-smokers showing increased learning from rewards in the PST test phase. Study 2

extended this with a larger sample, and the application of computational models of decision-making to the PST. This provided greater explanatory clarity regarding the higher-order cognitive processes involved in nicotine-dependence. Smokers and ex-smoker group membership was predicted similarly by reductions in drift-rate, which may be attributable to a general impairment in evidence accumulation observed in other externalizing disorders and psychopathologies. While the task-based features provided insight into relevant factors of decision-making, they did not outperform self-reported impulsivity in classifying smoker groups.

Chapter 6: Longitudinal changes in value-based decision-making as predictors of nicotine relapse

6. 1. Introduction

Although a large proportion of current smokers report a motivation to quit, only 3-5% of individuals remain successfully abstinent at 6-12 months following a quit attempt (Hughes et al., 2004). Considering various methods of estimating the frequency of smoking quit attempts, Chaiton et al. (2016) concluded that current smokers make an average of 30 quit attempts before successfully abstaining for at least one year.

A number of socio-economic and psychological factors have been identified in predicting successful nicotine abstinence. Lower socio-economic status groups show higher rates of nicotine use, lower treatment adherence, and in a study by Kotz and West (2009) the most deprived socio-economic group was half as likely to be successful in a quit attempt compared to the least deprived group (Hiscock et al., 2011; Reid et al., 2010). Additionally, level of nicotine dependence, number of cigarettes consumed daily, age of initial/first cigarette use, and the use of e-cigarettes or nicotine replacement therapy can also predict the success of quit attempts (Hajek et al., 2019; Hays et al., 2010). In a randomised controlled trial, Hajek et al. (2019) found that e-cigarettes were more effective than nicotine replacement therapy when combined with behavioural support (1-year abstinence rate of 18% versus 9.9%).

Changes in reward processing have been observed under conditions of nicotine abstinence versus satiation. Decreases in reward sensitivity have been observed in prolonged nicotine abstinence during a quit attempt (Hughes et al., 2017), however some studies report no change in reward sensitivity, an inconsistency that may be due to differences in experimental tasks and methods of analyses between studies. Some studies show electrophysiological changes associated with nicotine abstinence – for example, ex-smokers showed similar P300 ERPs to non-smokers compared with current smokers in a cue reactivity task. However, both former and current smokers showed diminished P300 on an auditory oddball task (Neuhaus et al., 2006). A recent study by Bu et al. (2019) found that EEG theta coherence in frontal-parietal

regions was predictive of cigarette craving, while others have shown that conflict-related theta may be associated with other forms of substance dependence (Harper et al., 2018a). Increases in medial-frontal theta power have been associated with increased conflict and uncertainty in decision-making tasks, and trial-to-trial increases in theta show a linear relationship with boundary separation in the DDM during high-conflict conditions in the test phase of the PST (Cavanagh et al., 2014). Despite this relationship, relatively few studies have investigated changes in electrophysiological markers of conflict-related cognitive control during nicotine abstinence.

In a within-sample study manipulating smoking status, Baker et al. (2018) showed that smoking status modulated learning rates during the probabilistic selection task, such that positive learning rates were enhanced following cigarette consumption, and reduced following abstinence. The reverse pattern was observed for negative learning rates, highlighting the utility of reinforcement learning models in characterising addiction states. In an application of an evidence accumulation model to a flanker task, Weigard et al. (2018) reported increased drift-rate variability (indicative of mind-wandering), and reductions in boundary separation among current smokers under conditions of abstinence. However, the majority of studies examining computational differences in decision-making among addiction groups are cross-sectional in design (Gueguen et al., 2020), and therefore it is unclear if such computational markers of decision-making can predict likelihood of abstinence versus relapse.

The current study sought to examine changes in value-based decision-making as predictors of time to nicotine relapse. In a longitudinal study design with biochemically verified nicotine abstinence, participants completed task-based and self-report measures during a baseline (T0) testing session immediately prior to a smoking quit attempt, and completed follow-up assessments at 1-week (T1) and 4-weeks (T2) post quit attempt if they remained abstinent. Feedback-related ERPs were compared at T0, T1, and T2. There were two aims: first,

an exploratory machine learning analysis with Cox proportional hazards models was used to predict time to nicotine relapse, testing the relative predictions of models with behavioural, computational, and EEG features. Second, the ERPs of those who remained abstinent were analysed: it was hypothesised that feedback-related ERPs would become attenuated during abstinence.

6. 2. Methods

6. 2. 1. Participants

Participants were recruited via community advertising (e.g., posters on campus, health centers, pharmacies), and via advertisements placed on social media focused upon the greater Dublin region (e.g., using Facebook). Participants received up to €10 in receipted travel expenses for each attendance as part of the study. The final sample consisted of 112 participants (71 female, 41 male; mean age 46.78, SD=11.66). CO and salivary cotinine levels were used to validate smoking abstinence versus relapse. The study was approved by the University College Dublin, and Trinity College Dublin, School of Psychology Ethics Committees.

6. 2. 2. Materials

6. 2. 2. 1. Probabilistic Selection Task

As described in Chapter 2, Section 2. 1. 1. The reward and punishment probabilities for stimuli in the PST training phase were randomised between study sessions at T0, T1, and T2.

6. 2. 2. 2. Barratt Impulsivity Scale

As described in Chapter 2, Section 2.2.

6. 2. 2. 3. Fagerstrom Test for Nicotine Dependence (FTND)

As described in Chapter 2, Section 2.3.

6. 2. 2. 4. Hooked on Nicotine Checklist (HONC)

The Hooked on Nicotine Checklist (Wellman, 2005) assesses lost autonomy over use of tobacco.

6. 2. 2. 5. Short Form Smoking Consequences Questionnaire (S-SCQ)

Smoking outcome expectancies were evaluated with the Short Form Smoking Consequences Questionnaire (S-SCQ; Myers, MacPherson, McCarthy, & Brown, 2003), with four subscales (negative consequences, positive reinforcement, negative reinforcement and appetite/weight controls).

6. 2. 2. 6. EEG

EEG data were recorded using the ActiveTwo Biosemi™ system in a soundproofed, darkened room from 70 electrodes (64 scalp electrodes) organised according to the 10–5 system (Oostenveld and Praamstra, 2001). Participants were seated in front of a cathode ray computer monitor with a screen resolution of 1024x768, and pixel refresh rate of 75 Hz. Participants were instructed to maintain their focus on the stimuli presented on the screen during the experiment. Activity related to eye movement was recorded bilaterally from approximately 2 cm below the eye and from the outer canthi. EEG data pre-processing was carried out using the EEGLAB toolbox (Delorme & Makeig, 2004; <http://sccn.ucsd.edu/eeglab>) in conjunction with the FASTER plug-in (Fully Automated Statistical Thresholding for EEG artefact Rejection; Nolan, Whelan, & Reilly, 2010, <http://sourceforge.net/projects/faster>). The data were bandpass filtered between 0.1 and 95 Hz, notch filtered at 50 Hz and average referenced across all scalp electrodes. A low-pass filter was applied at 30Hz. Data were subsequently epoched from 700 ms pre-feedback stimulus to 1000 ms post-stimulus. Artefactual (i.e., non-neural) independent components were identified and removed from the EEG data automatically using FASTER, as were epochs containing large artefacts (e.g., muscle twitches). Channels with poor signal quality were interpolated. The EEG data were then visually inspected to ensure good quality and that any

remaining noisy data were removed. The data were subsequently re-referenced to the average mastoids. Event-related potential data were baseline adjusted (-200 to 0 ms).

A spatial filter was applied to the data using the spherical spline algorithm to compute current source density/surface Laplacian estimates for surface potentials using the CSD toolbox in Matlab (Kayser & Tenke, 2006; <https://psychophysiology.cpmc.columbia.edu/software/csdttoolbox/>). ERP and time-frequency analyses were conducted on the transformed CSD data. Power spectral density estimates were calculated using Welch's method (pwelch function in MATLAB), expressed as relative band power, and theta band activity was averaged between 4-8 Hz from 150-500 ms post-feedback stimulus. Time-frequency and ERP regions of interest were expected to be at the FCz electrode (in line with Cavanagh et al., 2014; 2019). Analyses were conducted on the FCz channel, chosen in line with previous research examining the FRN and P3a, and studies examining EEG correlates during the PST (Cavanagh et al., 2011; 2019; West et al., 2018). Peak ERP amplitudes were extracted for each subject by identifying the peak amplitude Topoplots at each study time-point are presented in Appendix D (Section D. 1).

6. 2. 3. Procedure

Individuals interested in participating were initially contacted via a brief phone call to assess eligibility (e.g., participants were required to smoke at least 5-10 cigarettes per day; Appendix E). Eligible participants who provided informed consent were scheduled for three laboratory testing sessions; a baseline session 24 prior to their agreed smoking quit date (T0), 1-week (+ 1 day) post quit date (T1), and 4-weeks (+ 4 days) post quit date (T2). Participants also provided their contact details for random breath and saliva spot checks, which were performed once between the T0 session and T1 session, and twice between the T1 and T2 testing sessions. Self-report questionnaires and PST task-related EEG were recorded at each laboratory testing session, and were part of a larger battery of tasks (for further study details, see Lespine et al. 2020; Lespine et al., Under Review).

Behavioural data analysis

Trials with RTs >10 or <0.1 s were removed from the data (65 trials were removed, 0.04% of all trials).

Drift-diffusion Models

As described in Chapter 2, Section 2.3.

6. 3. Results

Of the initial sample (N=109), 35 participants were abstinent one week post nicotine quit attempt, and 21 participants were abstinent at T2. Follow-up questionnaires showed that 15 participants remained abstinent at 6 months post quit attempt.

6. 3. 1. Behavioural analysis

Mean (SD) values of the behavioural variables in the PST with-feedback and no-feedback phases at each time-point are presented in Table 6. 1.

Table 6. 1. Mean (SD) questionnaire and bioverification scores at each study time-point (T0 = baseline, T1 = Week-1, T2 = Week=4). CO reading = Carbon Monoxide reading, DASS = Depression, Anxiety, Stress Scale, BIS = Barratt Impulsiveness Scale, S-SCQ = Short Smoking Consequences Questionnaire, HONC = Hooked on Nicotine Checklist, FTND = Fagerstrom Test for Nicotine Dependence.

	<i>T0</i>	<i>T1</i>	<i>T2</i>
<i>CO Reading</i>	19.7 (10.5)	2.2 (1.3)	2.3 (1.0)
<i>Cotinine</i>	4.0 (0.6)	0.9 (1.1)	0.7 (0.9)
<i>DASS Stress</i>	15.6 (11.1)	13.3 (10.9)	11.3 (9.8)
<i>DASS Anxiety</i>	10.1 (9.6)	4.9 (6.1)	4.4 (5.96)
<i>DASS Depression</i>	11.8 (10.9)	6.1 (6.2)	10.1 (11.7)
<i>BIS Attentional</i>	16.8 (3.7)		
<i>BIS Motor</i>	23.7 (2.3)		
<i>BIS Non-planning</i>	26.5 (5.5)		
<i>S-SCQ Neg. consequences</i>	34.1 (4.4)		
<i>S-SCQ Pos. reinforcement</i>	22.1 (12.4)		

<i>S-SCQ Neg. reinforcement</i>	39.2 (14.1)
<i>S-SCQ Appetite/Weight Control</i>	22.1 (11.6)
<i>AUDIT Total</i>	8.05 (6.2)
<i>AUDIT: Dependency</i>	1.14 (1.9)
<i>AUDIT: Hazardous Consumption</i>	4.69 (2.6)
<i>AUDIT: Alcohol related Harm</i>	2.22 (3.11)
<i>AUDIT: Alcohol Problems</i>	3.35 (4.46)
<i>DAST Total</i>	2.57 (2.8)
<i>Age Start Smoking</i>	16.9 (5.7)
<i>HONC Total</i>	8.3 (1.6)
<i>FTND Total</i>	2.4 (1.6)

6.3.1.1. PST Behavioural summary

PST Test phase performance was correlated between study time-points: mean Approach AC correlated between baseline and one-week ($r=0.47$, $p = .006$), and one-month ($r = .58$, $p = .005$; see Table D.1, Appendix D). This was also the case between Approach BD at baseline and one-week ($r = 0.5$, $p = .003$), and one-month ($r=.56$, $p = .008$). Bayesian paired-samples t-tests showed significant differences in RTs between baseline (T0) and week-one (T1) for total RT in the PST training phase (Mean T0 = 1.17, SD = 0.44), AB trials (T0 = 0.99, SD = 0.34), and CD trials (T0 = 1.16, SD = 0.5), with faster RTs at T1. Faster RTs were also observed at T1 for the Lose-Lose condition in the PST test phase (T0 = 1.41, SD = 0.51).

Table 6. 2. Mean performance (percentage choice accuracy) and reaction times (RT) in the Probabilistic Selection Task training and test phases.

<i>Phase</i>	<i>T0 (n=109)</i>	<i>T1 (n=31)</i>	<i>T1 (n=21)</i>	<i>BF10 (T0 vs. T1)</i>
Training				
Accuracy (% Choice)				
<i>AB</i>	82.49 (16.28)	80.65 (22.41)	83.52 (19.57)	0.28
<i>CD</i>	79.19 (15.93)	78.38 (20.3)	76.12 (18.29)	0.18
<i>EF</i>	54.16 (19.13)	50.29 (18.09)	50.68 (17.20)	0.19
RT (s)				
<i>Total</i>	1.25 (0.94)	0.98 (0.48)	1.10 (1.01)	21.81
<i>AB</i>	1.10 (0.8)	0.82 (0.66)	0.91 (0.71)	195.08
<i>CD</i>	1.25 (0.99)	0.96 (0.8)	1.09 (1.04)	13.62
<i>EF</i>	1.41 (0.99)	1.15 (1.04)	1.28 (1.16)	3.55
Test				
Accuracy (% Choice)				
<i>Approach AC</i>	70.02 (16.58)	73.44 (20.86)	76.64 (19.68)	0.26
<i>Approach BD</i>	52.04 (16.13)	64.25 (19.19)	63.12 (15.49)	7.513
<i>Total RT</i>	1.14 (0.49)	1.0 (0.5)	0.90 (0.59)	3.27
<i>Win-Win RT</i>	1.03 (0.46)	0.94 (0.45)	0.83 (0.52)	0.39
<i>Win-Lose RT</i>	1.01 (0.46)	0.91 (0.48)	0.82 (0.58)	0.64
<i>Lose-Lose RT</i>	1.42 (0.65)	1.22 (0.59)	1.11 (0.73)	10.14

6. 3. 2. ERP analysis

EEG Data Quality

Four datasets were removed due to poor EEG data quality (labelled as missing data in Cox regression), and data from three subjects were removed due to missing follow-up information regarding nicotine relapse. The final sample consisted of 109 subjects at baseline, 29 subjects at T1, and 22 subjects at T2.

Following pre-processing, the mean number of positive RPE trials was 60.86 (SD=15.77) and mean number of negative RPE trials was 40.91 (SD=9.15). ERPs were compared at the FCz electrode.

Whole Group

ERPs




Figure 6. 1. Whole group ERPs for positive (RPE+) and negative (RPE-) feedback trials, locked to the feedback stimulus onset (0).

ERP Group Comparisons

Two group comparisons were made: (i) Baseline (T0) ERPs for those who relapsed within the first week of their quit attempt ($n = 74$), versus those who were successfully abstinent at one-week ($n = 35$; Fig 1), and (ii) Baseline (T0) versus T1 ERPs for abstinent participants ($N = 33$; Fig 2).

T0 Comparisons

There were no significant differences in the number of positive and negative RPE trials analysed for T0 participants who successfully abstained versus those who relapsed at T1 (p 's > .45). Mean

baseline ERPs for positive (RPE +) and negative (RPE -) feedback trials between those who remained abstinent at T1 versus those who relapsed within the first week of their quit attempt are presented in Figure 6. 2.





Figure 6. 2. Group comparisons of feedback-locked ERPs at T0 between those who relapsed within T1 and those who were abstinent at T1.

T0 vs. T1 Comparisons

There were no significant differences in the number of positive and negative RPE trials for the successfully abstinent participants between T0 and T1 (p 's > .22). ERPs between T0 and T1 are presented in Figure 6. 3.




Figure 6. 3. Group comparisons of feedback-locked ERPs at T0 versus T1.

Feedback-related ERPs for successful abstainers (n=19) across the three study time-points are presented in Figure 6. 4.

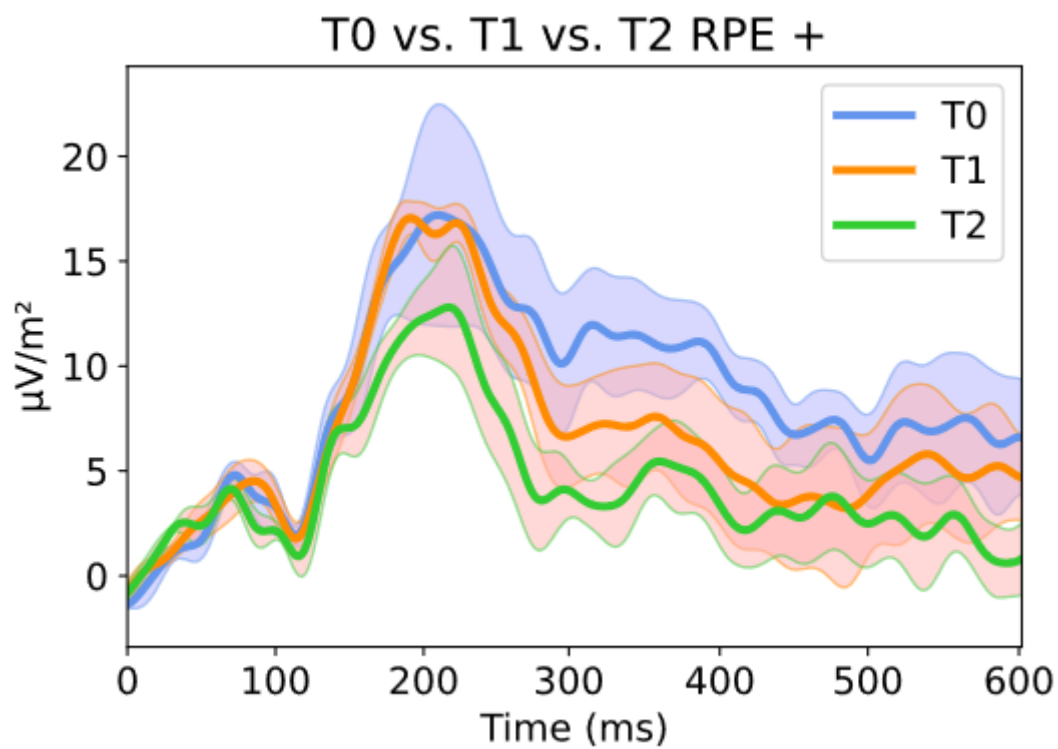



Figure 6. 4. Feedback-locked ERPs for positive (RPE +) and negative (RPE -) feedback trials at T0, T1, and T2.

PST Test Phase ERPs

Group differences in the PST test phase were explored by comparing ERPs between conflict conditions (Lose-Lose, Win-Lose, Win-Win; Figure 6. 5) and between ERPs locked to responses (Figure 6. 6).




Figure 6. 5. Comparison between Win-Win, Win-Loss, and Loss-Loss conditions for the whole sample.




Figure 6. 6. Comparison of mean ERPs time-locked to the response trigger between abstinent participants at T0 versus T1.

6. 3. 3. Computational Modelling Results

6. 3. 3. 1. Behavioural vs. Neural data RLHDDM Models

Feedback-locked EEG theta-band power (4-8 Hz) from the FCz electrode was extracted during the PST training phase as an index of mPFC activity. To examine the relationship between theta and reaction time, a single-trial multiple regression was conducted for each subject, the mean standardised beta was 0.015 (.05). A comparison of single-trial betas from a regression between feedback type and RT on the following trial showed a significant difference between error/punishment trials (0.024, SD = 0.12) and reward/positive feedback trials (-0.005, SD = 0.08), $W = 11534.0$, $p = 0.0031$. This was the case for the baseline (T0) participants ($n = 112$) and

a similar trend was observed at T1 ($n = 33$) (W 's = 5178.0, 428.0 and p 's = 0.007, 0.068 respectively).

6.3.3.2. Drift-diffusion Computational Models of PST

A number of HDDM and RLHDDMs were compared (see Appendix D). An RLHDDM model with theta as a trial-by-trial regressor on the boundary separation parameter, and feedback type as an interaction term was entered into the ML analysis. The HDDM was applied to behavioural data, and was the same as that applied in previous chapters.

No significant paired group differences in mean model parameters between T0 and T1 were observed, and boundary separation regression parameters did not differ between positive and negative feedback at any time-point (p 's > 0.32).

Table 6. 3. Comparison of mean RLHDDM and HDDM parameter estimates between time-points, with Bayesian paired sample t-test between T0 and T1 (BF10).

<i>Parameter</i>	<i>T0 (n=109)</i>	<i>T1 (n=31)</i>	<i>T2 (n=21)</i>	<i>Sig. Test (BF10)</i>
<i>RLHDDM</i>				
<i>Boundary Separation: Intercept</i>	1.76 (0.4)	1.58 (0.48)	1.6 (0.53)	0.36
<i>Boundary Separation: theta, pos. feedback</i>	0.04 (.05)	0.023 (.05)	0.04 (.05)	0.34
<i>Boundary Separation: theta, neg. feedback</i>	0.03 (.05)	0.02 (.05)	0.02 (.05)	0.31
<i>Learning Rate</i>	0.97 (0.4)	0.79 (0.39)	0.85 (0.35)	0.19
<i>Drift-rate</i>	2.14 (0.69)	2.09 (1.54)	2.07 (0.97)	0.72
<i>Non-decision Time</i>	0.32 (0.15)	0.25 (0.11)	0.25 (0.11)	0.22
<i>HDDM</i>				
<i>Boundary Separation</i>	1.89 (0.47)	1.76 (0.52)	1.65 (0.56)	3.29
<i>Drift-rate (Loss-Loss)</i>	0.11 (0.45)	0.25 (0.51)	0.14 (0.37)	0.19
<i>Drift-rate (Win-Loss)</i>	1.26 (0.98)	1.12 (1.11)	1.26 (0.96)	0.27
<i>Drift-rate (Win-Win)</i>	0.89 (0.88)	0.80 (1.10)	0.92 (0.81)	0.23
<i>Non-decision Time</i>	0.35 (0.15)	0.32 (0.14)	0.28 (0.12)	0.22

6. 3. 4. Cox Regression

The PST Summary Model mean concordance index was 0.495 (.03) with a LL of 46.31 (0.2), and was significantly lower than the demographic model was 0.515 (.03) with a LL of 46.26; $t(198) = 4.89$, $p < .001$. The demographic model outperformed the PST Summary model in 66% of iterations.

Table 6. 4. Mean choice frequency and beta values from the PST Summary model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Null</i>	<i>Mean Beta - Null</i>
<i>Age</i>	4.11	0.040	4.19	0.048
<i>Gender</i>	2	-0.011	2.66	-0.025
<i>Approach BD</i>	1.55	-0.006	2.21	0.016
<i>Approach AC</i>	1.43	-0.013	3.64	-0.034
<i>Approach AC (T0)</i>	1.37	0.012	3.59	-0.033
<i>Approach BD (T0)</i>	0.95	-0.001	1.54	0.008

The Computational Model mean concordance index was 0.54 (.03) and LL = 46.29 (0.4), and was significantly higher than the demographic model C-index was 0.52 (.02) and LL = 46.27 (0.3); $t(198) = 6.13$, $p < .001$. Baseline boundary separation from the PST test phase, baseline non-decision time from both PST phases, and changes in non-decision times during the PST training phase were significant predictors.

Table 6. 5. Mean choice frequency and beta values from the Computational model.

<i>Feature</i>	<i>Mean choice frequency - Test</i>	<i>Mean Beta - Test</i>	<i>Mean choice frequency - Null</i>	<i>Mean Beta - Null</i>
<i>HDDM: Boundary Separation – T0</i>	4.43	-0.131	4.57	0.0713
<i>HDDM: Non-decision Time – T0</i>	4.29	0.082	0.46	0.002
<i>RLHDDM: Non-decision Time – T0</i>	4.26	0.067	0.42	0.0004
<i>RLHDDM: Non-decision Time</i>	3.57	0.042	0.15	0.0004
<i>RLHDDM: Drift-rate</i>	3.1	-0.036	1.79	-0.016
<i>Age</i>	2.49	0.021	2.3	0.022
<i>HDDM: Drift-rate (Win-Lose) – T0</i>	2.19	0.036	0.22	0.0001

<i>HDDM: Drift-rate (Lose-Lose)</i>	2.17	-0.016	1.29	0.011
<i>HDDM: Drift-rate (Lose-Lose) – T0</i>	1.54	-0.016	0.42	-0.003
<i>Gender</i>	1.51	-0.014	0.8	-0.006
<i>HDDM: Non-decision Time</i>	1.47	0.005	0.31	-0.001
<i>RLHDDM: Boundary Separation (Pos. feedback) – T0</i>	1.2	0.011	0.25	0.0009
<i>RLHDDM: Boundary Separation (Intercept)</i>	1.19	0.014	0.43	0.001
<i>HDDM: Drift-rate (Win-Lose)</i>	0.79	0.002	0.42	-0.002
<i>HDDM: Drift-rate (Win-Win)</i>	0.74	-0.006	0.26	-0.0005
<i>RLHDDM: Learning rate – T0</i>	0.66	0.006	0.15	0.0006
<i>HDDM: Drift-rate (Win-Win) – T0</i>	0.61	-0.006	3.25	-0.045
<i>RLHDDM: Boundary Separation (Pos. feedback)</i>	0.57	0.002	0.9	-0.007
<i>HDDM: Boundary Separation</i>	0.52	0.007	2.84	0.023
<i>RLHDDM: Learning Rate</i>	0.52	-0.002	0.44	0.003
<i>RLHDDM: Drift-rate – T0</i>	0.49	0.001	0.81	-0.003
<i>RLHDDM: Boundary Separation (Neg. feedback)</i>	0.48	0.0001	1.72	-0.0127
<i>RLHDDM: Boundary Separation (Neg. feedback) – T0</i>	0.47	0.001	1.64	-0.013
<i>RLHDDM: Boundary Separation (Intercept) – T0</i>	0.38	0.0001	0.26	0.001

The C-index for the model with mean computational parameters and peak ERP amplitudes (Mean C-index = 0.53, SD = .03; LL = 46.36, SD = 0.4) outperformed the demographic model (with age and gender as fixed covariates, and random data; C-index = 0.51 (LL= 46.41); $t(198) = 6.77$, $p < .001$). Significant features were the same as the Computational Model (i.e., boundary separation and non-decision times), however the ERP features were not frequently selected.

Table 6. 7. Mean choice frequency and beta values from the Computational and EEG model.

Feature	Mean choice frequency - test	Mean beta - test	Mean choice frequency - Demographic	Mean beta - Demographic
<i>HDDM: Boundary Separation – T0</i>	4.17	-0.123	0.85	-0.004
<i>HDDM: Non-Decision Time –T0</i>	4.15	0.074	0.52	-0.002
<i>RLHDDM: Non-Decision Time – T0</i>	4.1	0.065	0.23	0.001
<i>RLHDDM: Non-Decision Time</i>	3.5	0.039	1.57	-0.011
<i>RLHDDM: Drift-Rate</i>	2.87	-0.036	1.11	-0.007
<i>Age</i>	2.28	0.020	2.38	0.019
<i>PSWFB: ERPS (Pos. feedback) – T0</i>	2.2	-0.026	0.39	0.002
<i>HDDM: Drift-Rate (Win-Lose) – T0</i>	1.99	0.030	1.57	0.009
<i>HDDM: Drift-Rate (Lose-Lose)</i>	1.88	-0.012	1.01	0.006
<i>HDDM: Drift-Rate (Lose-Lose) – T0</i>	1.39	-0.014	1.69	0.015
<i>Gender</i>	1.33	-0.013	0.71	-0.005
<i>HDDM: Non-Decision Time</i>	1.2	0.005	1.69	0.013
<i>RLHDDM: Boundary Separation (Intercept)</i>	1.06	0.015	0.55	0.002
<i>RLHDDM: Boundary Separation (Pos. Feedback) – T0</i>	1.05	0.011	0.99	0.010
<i>PSTNFB: ERP - Baseline</i>	0.86	-0.010	1.45	-0.009
<i>RLHDDM: Learning Rate – T0</i>	0.81	0.007	0.31	0.0004
<i>HDDM: Drift-Rate (Win-Win)</i>	0.69	-0.004	0.65	0.004
<i>PSWFB: ERPS (Neg. feedback)</i>	0.58	-0.001	0.91	-0.005
<i>HDDM: Drift-Rate (Win-Lose)</i>	0.56	0.0001	0.59	-0.003
<i>RLHDDM: Learning Rate</i>	0.54	-0.002	2.6	0.030
<i>HDDM: Boundary Separation</i>	0.52	0.003	0.41	-0.0006
<i>PSWFB: ERPS (Neg. feedback) – T0</i>	0.51	0.004	1.98	-0.016

<i>HDDM: Drift-Rate (Win-Win) – T0</i>	0.48	-0.006	0.84	-0.007
<i>RLHDDM: Boundary Separation (Pos. Feedback)</i>	0.46	0.002	1.61	0.012
<i>RLHDDM: Boundary Separation (Neg. Feedback)</i>	0.44	-0.0007	4.64	0.093
<i>PSWFB: ERPS (Pos. feedback)</i>	0.42	0.005	3.27	-0.041
<i>RLHDDM: Boundary Separation (Neg. Feedback) – T0</i>	0.41	-0.0007	0.29	-0.0001
<i>PSTNFB: ERP</i>	0.36	0.0003	0.41	0.002
<i>RLHDDM: Drift-Rate – T0</i>	0.35	1.75679E-05	0.41	-0.0009
<i>RLHDDM: Boundary Separation (Intercept) – T0</i>	0.32	-0.00	2.02	0.020

6. 3. 5. Correlations between ML features

Correlations between ML features (Table 6. 8) at T0 showed significant relationships within the self-report measure, and correlations between computational model parameters. Correlations between computational parameters at each study time-point are presented in Appendix D.

Feature	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1. BIS Attention	—																	
2. BIS Motor	0.421***	—																
3. BIS Non-planning	0.565***	0.525** *	—															
4. RLHDDM: <i>a</i> Intercept	0.145	-0.066	-0.098	—														
5. RLHDDM: <i>a</i> (RPE+)	-0.055	-0.029	-0.081	- 0.304**	—													
6. RLHDDM: <i>a</i> (RPE-)	-0.141	-0.136	-0.161	-0.259*	0.861***	—												
7. RLHDDM: Learning Rate	-0.134	-0.153	-0.087	-0.130	-0.015	0.032	—											
8. RLHDDM: <i>t</i>	-0.005	-0.024	0.007	0.069	-0.142	-0.072	0.261*	—										
9. RLHDDM: <i>v</i>	-0.089	-0.008	-0.011	-0.197	-0.097	-0.021	0.307**	0.133	—									
10. HDDM <i>a</i>	-0.174	-0.216	-0.019	0.413**	0.123	0.173	0.235	0.209	0.074	—								
11. HDDM <i>vLL</i>	-0.008	0.003	0.078	-0.186	-0.048	-0.083	0.183	0.021	0.209	- 0.006	—							
12. HDDM <i>vWL</i>	-0.033	-0.014	0.020	-0.053	-0.166	-0.097	0.291*	0.224	0.159	0.056	0.432** *	—						
13. HDDM <i>vWW</i>	-0.064	-0.021	-0.068	-0.138	-0.221	-0.147	0.335**	0.235	0.349 *	- 0.108	0.376** *	0.376** *	—					
14. HDDM: <i>t</i>	-0.0195	0.033	0.057	0.105	0.157	0.089	0.100	0.288*	-0.047	0.400	-0.244	-0.224	-0.268*	—				
15. ERP (RPE +)	0.223	0.065	0.139	0.176	0.045	-0.004	-0.061	-0.069	-0.204	0.059	-0.251*	-0.251	-0.211	0.001	—			
16. ERP (RPE-)	0.116	-0.024	0.073	0.189	0.085	0.049	0.118	0.0001	-0.103	0.154	-0.215	-0.215	-0.129	0.138	0.674** *	—		
17. ERP (Test phase response)	-0.128	-0.051	0.139	-0.075	-0.055	-0.105	-0.005	-0.014	-0.062	- 0.143	0.054	0.054	0.036	-0.031	-0.098	-0.161	—	

18. Approach AC	-0.059	-0.056	-0.066	-0.056	-0.215	-0.104	0.363**	0.276*	0.321*	0.044	0.291*	0.291*	0.889**	-0.172	-0.223	-0.087	-0.047	
19. Approach BD	-0.009	-0.019	0.047	-0.037	0.032	0.014	0.265*	0.097	0.068	0.184	0.727**	0.727**	0.239	-0.112	-0.107	0.026	-0.033	0.336*

*p < .05, ** p < .01, *** p < .0001

Figure 6. 8. Partial Spearman's correlations between machine learning features at baseline, with age as a covariate. Significant correlations are highlighted between (i) self-report measures (orange), (ii) task-related parameters and variables (blue), (iii) task-related and self-report measures (green)

6. 4. Discussion

The present study sought to examine EEG-informed computational parameters of decision-making as predictors of time to nicotine relapse in a longitudinal smoking cessation study. Behaviourally, participants at T1 (i.e., Week-1) of successful abstinence showed reduced learning from negative feedback in the PST test phase (i.e., higher 'Approach BD') compared with baseline. Attenuated ERPs were observed across study time-points in response to positive and negative feedback in the training phase.

Cue-related P3 amplitudes have been associated with violations of expectation, such that higher amplitudes are elicited from unexpected relative to expected reward and punishment feedback (Donaldson et al., 2016). Attenuated FRN amplitudes have been related to increases in dopamine activity (Holroyd & Coles, 2002). Although a number of studies have investigated acute effects of nicotine abstinence on EEG correlates of performance monitoring, few have examined the effects of longer term abstinence during smoking cessation attempt. Schliez, Hawk, and Rosch (2013) showed reduced error-related negativity on a Flanker task during acute nicotine abstinence, however no effect on N200 and P300 components. Acute nicotine abstinence (12 hours) was also shown to reduce P3b amplitudes during an oddball task, and selectively reduce P3a amplitudes among current smokers reporting lower cognitive control (Evans et al., 2013). The ERP results presented here suggest that P3 amplitudes maybe attenuated in longer-term (i.e., 1 week) abstinence during smoking cessation, whereas FRN amplitudes during negative feedback trials in the PST are more pronounced at Week 1 compared with baseline..

Increased FRN in response to negative feedback on the PST was found among patients with major depressive disorder compared to controls (Cavanagh et al., 2011), and was associated with increased learning from negative feedback in the test phase. Here, no significant correlation was observed between Approach BD and mean ERP (RPE-) during the training phase,

this indicates that neural processing of feedback in the training phase did not correspond with value-based decision-making during the test phase. It remains possible that other methods of indexing feedback processing (e.g., by taking the difference between the P2/P3 and FRN in the training phase) may show a relationship with behavioural performance in the PST test phase. However, there was a trend towards a negative relationship between mean feedback-related amplitudes in the training phase and drift-rates in the test phase. This suggests that smaller drift-rates (i.e., faster accumulation of evidence) may be associated with increased feedback-related ERP amplitudes during the training phase.

RLHDDMs with mid-frontal theta as a single-trial regressor were fit to the PST training phase, and a HDDM was fit to the behavioural data of the test phase. No group differences in computational parameter estimates were observed between baseline and week-one, however the boundary separation parameter from the test phase showed reductions across study time-points that were close to significance. The influence of feedback-related trial-by-trial theta on boundary separation in the PST training phase was modelled as a regressor in the RLHDDM, this tested the hypothesis that trial-by-trial theta shows a linear relationship with boundary separation during the training phase that interacts with feedback type. However, no significant difference in the boundary separation parameter between positive and negative feedback conditions was observed.

Penalised cox regression models with Computational model parameter estimates, and mean ERP amplitudes as features outperformed demographic models with age and gender at predicting time to nicotine relapse. On the other hand, the PST Summary model did not outperform demographic models in predicting time to nicotine relapse. The Computational models showed that smaller baseline boundary separation, and larger non-decision times were the strongest features predicting longer time to relapse. These features were followed by

changes in non-decision times and drift-rates from the PST training phase across study time-points, age, and peak ERP amplitudes in response to positive feedback.

Correlations between the ML features at T0 showed significant relationships between the computational parameter and PST summary scores. Positive relationships were found between mean boundary separation regression parameters for positive and negative feedback, and mean ERP amplitudes for positive and negative feedback also positively correlated. These results suggest that feedback type did not bidirectionally affect these predictors.

A limitation of the current chapter is the high relapse rate at T1 and T2, which limits the inferences that can be made regarding group differences in self-report, task-based, and EEG measures. Of the initial sample, 67% had relapsed within the first week – future research may consider more dense experimental assessments within the early stage of smoking cessation to identify predictors of early relapse in this group. Furthermore, it is possible that changes in PST performance across study time-points may have resulted from familiarity with the experimental task. However, this was partially mitigated by randomising the stimuli reward probabilities in the PST between study sessions, and similar repeated designs have been reported elsewhere in the literature (e.g., Baker et al., 2018).

In summary, these findings suggest subtle behavioural and electrophysiological reward-related changes during smoking cessation that may be predictive of time to relapse. This highlights the utility of repeated measures of decision-making during the early abstinence period in smoking cessation. The use of cognitive models such as the DDM may provide more precise information regarding the latent psychological processes involved in successful abstinence, such as changes in boundary separation and non-decision times.

Chapter 7: General Discussion

This chapter will provide synthesis of the main thesis findings, novel contributions, and future directions for research. The current programme of research aimed to provide a comprehensive investigation of value-based decision-making in substance dependence by comparing behavioural, computational, and electrophysiological features to predict various substance misuse and dependence phenotypes.

Chapter 1 provided an overview of decision-making research in the field of substance-dependence, recent advances in computational psychiatry, and clinical applications of machine learning. This highlighted that although substance-dependence is associated with reward-related decision-making impairments, methodological heterogeneity regarding the types of experimental tasks and populations studied has led to inconsistencies in the literature. For example, it is not clear how successful abstinence affects decision-making, and whether this differs according to type of substance use and treatment (Stewart et al., 2019). Computational psychiatry, and data-driven machine learning methods are well-suited to addressing these issues, by identifying task-general latent factors that may predict psychopathology trans-diagnostically, and may offer a means to more precise treatment interventions (Gillan & Whelan, 2017; Yarkoni & Westfall, 2017). The drift-diffusion model is a prominent decision-making model that incorporates latent features (e.g., evidence accumulation, response caution) thought to underlie neurocognitive deficits in psychopathology (Sripada & Weigard, 2021). In this thesis, I fit combined reinforcement learning and drift-diffusion models to a reward learning task that has been extensively described in the literature, and which is thought to index fluctuations in dopaminergic functioning and show sensitivity to various clinical disorders (Frank et al., 2004).

In Chapter 2, the experimental task, questionnaire measure, computational modelling, and penalised regression approach were introduced. Chapter 2 described how, in the empirical chapters, RL(HDDM)s were fit to trial-by-trial PST data, and the parameters estimated from these were entered into penalised linear and logistic regression models with out-of-sample

predictions of various substance phenotypes. Chapter 3 sought to predict hazardous alcohol-use risk. Chapter 4 sought to predict length of abstinence in a sample formerly dependent on heroin. Chapter 5 sought to classify individuals based on their smoking group category (non-smokers, current smokers, ex-smokers, and vapers). Chapter 6 aimed to assess longitudinal changes in features among a sample of individuals undergoing a smoking cessation attempt. Sample-general alterations in latent cognitive constructs were revealed, and the relative predictive utility of self-report versus task-based features was compared.

7. 1. Summary of Results

The direction of significant relationships between computational parameter features and outcome variables from each empirical chapter are outlined in Table 7. 1. In a sample of adults with varying levels of alcohol consumption, the results from Chapter 3 showed that high alcohol misuse risk was predicted by increased negative learning rate, smaller boundary separation, larger drift-rates, smaller non-decision times, and higher impulsivity on all of the BIS 2nd order sub-scales. Robust partial correlations accounting for variability in computational model parameter estimates and controlling for the effect of age showed similar relationships with the total AUDIT score and its sub-scales.

Chapter 4 found that smaller drift-rates in the PST test phase, wider response boundaries, larger non-decision times, and greater learning rates in the PST training phase predicted longer length of abstinence from opioids. Higher self-reported sensation-seeking and motor impulsivity were also significant predictors. These findings were supported by robust negative correlations between drift-rates and length of abstinence, and positive correlations between boundary separation and length of abstinence. The variance in correlations with positive and negative learning rate parameter estimates suggested a less robust relationship with length of abstinence.

In Chapter 5, current smoker, ex-smoker, vaper, and non-smoker group membership was predicted by varying features. Relative to non-smokers, the current smoker group was

classified by smaller boundary separation and reduced positive learning rates in the PST training phase, larger non-decision times (both PST phases), smaller drift-rates for the high conflict Win-Win condition in the PST test phase, and higher self-reported non-planning impulsivity. Reduced drift-rates for the Loss-Loss condition, and increased drift-rates for Win-Win in the PST test phase, in addition to larger non-decision times, increased motor impulsivity, and decreased non-planning impulsivity were predictive of ex-smoker group membership. Of the computational parameters, the vaper group was classified by larger non-decision times.

Longitudinal changes in value-based decision-making during smoking cessation were investigated in Chapter 6, with self-report and task-based measurements recorded immediately prior to a smoking cessation attempt (T0), and at one-week (T1) and four-weeks (T2) following successful nicotine abstinence. Behaviourally, reduced learning from negative feedback during the PST test phase was observed at T1 compared with T0, indicating that abstinence resulted in an insensitivity to learning from punishers. Faster reaction times were also observed on a number of PST conditions at T1. Group differences in positive and negative feedback-related ERPs were observed, with attenuated activity observed at T1 and T2 compared with T0. However, differences in task performance across study time-points may be influenced by participants' familiarity with the task. Mean peak ERP amplitudes and computational model parameters were entered as time-dependent features into a penalised survival regression model to predict time to nicotine relapse, and outperformed demographic models. Baseline boundary separation and non-decision time were the most frequently selected features. These were followed by changes in non-decision time, drift-rate from the training phase, age, and mean baseline ERP amplitudes during reward trials in the PST training phase.

Correlations between machine learning model features in each empirical chapter revealed relationships between task-based features, between self-reported impulsivity, and between task-based and self-report features. Across all studies, the boundary separation parameter from the PST training phase showed a significant positive relationship with boundary

separation from the test phase, this was also observed for the non-decision time parameter. This finding indicates that individual differences in DDM parameters are stable across the training and test phase of the PST. In all studies, at least one of the drift-rate parameters correlated with PST test phase summary scores. Larger drift-rates were generally associated with greater Approach AC and Approach BD choice selections in the test phase, suggesting that speed of evidence accumulation is related to behavioural performance on the PST for both accurate (i.e., Approach AC) and inaccurate (i.e., Approach BD) choices. Correlations between task and self-report measures varied between studies- however boundary separation correlated significantly with motor impulsivity in Chapters 4 and 5.

	<i>RLHDDM: Boundary Separation</i>	<i>RLHDDM: Non- decision time</i>	<i>RLHDDM: Drift-rate</i>	<i>RLHDDM: Positive Learning Rate</i>	<i>RLHDDM: Negative Learning Rate</i>	<i>HDDM: Boundary Separation</i>	<i>HDDM: Non- decision time</i>	<i>HDDM: Drift-rate (Lose- Lose)</i>	<i>HDDM: Drift-rate (Win- Lose)</i>	<i>HDDM: Drift-rate (Win- Win)</i>
<i>Ch. 3. Higher alcohol-misuse risk</i>	↓	↓	↑	↓	↑	↓		↓	↑	↑
<i>Ch. 4. Shorter length opioid abstinence</i>	↓	↓		↓	↓	↑		↑	↑	↑
<i>Ch. 5. Smoker vs. Non-smoker</i>	↓	↑		↓			↑			↓
<i>Ch. 5. Ex-smoker (vs. Non-smoker, Current Smoker, Vaper)</i>		↑								
<i>Ch. 5. Vaper (vs. Non-smoker, Current Smoker, Vaper)</i>							↑			
<i>Ch. 6. Longer time to nicotine relapse</i>		↑				↓	↑			

Table 7. 1. Summary of thesis results from Combined Models with computational parameters. Arrows indicate direction of relationship between parameters and outcome variables from each empirical chapter. RLHDDM = Reinforcement learning hierarchical drift-diffusion model. HDDM = Hierarchical drift-diffusion model.

7. 2. Relationship with existing literature

Boundary Separation

Across all studies, the boundary separation parameter was a significant predictor of substance-use. This parameter indexes level of caution in decision-making, whereby wider bounds require more evidence accumulation before a response is reached. An individual with an emphasis on speed relative to accuracy is hypothesised to set narrow decision bounds, requiring less evidence to make choices, which has been termed ‘impulsive information processing’ (Metin et al., 2013). In the current set of studies, lower boundary separation significantly predicted (i) higher alcohol misuse risk, (ii) shorter abstinence from opioids, (iii) smoker versus non-smoker group membership, and (iv) a trend towards longer time to nicotine relapse.

These findings suggest that reduced boundary separation may present a general risk-factor for substance-use and severity, which seems to correspond conceptually with dual-process accounts of decision-making and substance dependence. That is, reduced boundary separation may reflect automatic, immediate, and impulsive behaviour relative to goal-directed and planned behaviour. Empirically, substance-dependent groups and university students with substance misuse behaviour tend to show dysregulation of inhibitory control on Go/No-Go and Stop-Signal tasks (Byrne & Worthy, 2019; Smith, Mattick, Jamadar, & Iredale, 2014). Groups with externalising disorders, such as ADHD, also demonstrate difficulties inhibiting responses and regulating speed-accuracy trade-offs in decision-making. Pederson et al. (2017) found that this resulted in reduced boundary separation during instrumental learning on the PST. I extend this finding to suggest that reductions in boundary separation also show a general relationship with substance-dependence that is sensitive to length of MMT treatment in opioid dependence. Interestingly, Lawlor et al. (2019) found increased boundary separation in a probabilistic reward task among patients diagnosed with depression compared with controls, therefore it is possible

that dissociable relationships with boundary separation exist across symptom dimensions. For example, externalising psychopathologies and disorders of compulsivity may be associated with reduced response caution, whereas disorders with symptoms of rumination and depressed mood may be associated with increased response caution.

Although individual differences in response boundary have been shown to discriminate between various clinical groups (e.g., Powell et al., 2019), Hedge et al., (2020) reported low correlations between boundary separation and self-reported impulsivity on the UPP-S questionnaire. In a meta-analysis of 19 studies, Spearman's rhos ranged from -0.04 to -0.02, indicating no relationship despite theoretical and conceptual similarities between boundary separation and self-reported impulsivity. I also report small/non-significant correlations between boundary separation and the SUPPS among former opioid users in the Supplementary Analysis (Appendix B) of Chapter 4. However, elsewhere in Chapter 4 (and Chapter 5) mean boundary separation estimates correlated negatively with the motor sub-scale of the BIS- consistent with the idea that smaller decision bounds are associated with increased impulsive information processing. This was further confirmed in robust correlations between the BIS motor sub-scale and samples of boundary separation estimates (Appendix B). On the other hand, no relationship between motor impulsivity and boundary separation was observed in Chapter 3, or in correlations between ML features at baseline in Chapter 6. The discrepancies between these results may be due to differences in sample characteristics, for example higher mean age in Chapters 4 and 6 compared to Chapters 3 and 5 may have influenced DDM parameters. Therefore, although reduced boundary separation was a significant predictor of alcohol misuse, it did not show a relationship with self-reported impulsivity in all studies.

Drift-rate

Drift-rate, or evidence accumulation, was also a significant predictor across studies. It has been suggested that evidence accumulation efficiency is an individual difference dimension

that may explain deficits of neuro-cognition in a host of clinical disorders (Weigard & Sripada, 2020). To explore the role of conflict in drift-rates during the PST test phase, separate parameters were estimated for trials that included ‘Lose-Lose’ (high conflict), ‘Win-Lose’ (low conflict), or ‘Win-Win’ (high conflict) stimulus pairs. Value differences have been shown to directly influence drift-rates during the PST (Cavanagh et al., 2014). Longer reaction times are typically observed on trials with two stimuli with similar reinforcement probabilities in the PST test phase, and shorter RTs on trials with divergent probabilities. Here, the drift-rate conflict conditions differentially predicted substance use – in Chapters 3 and 4, slower drift-rates across all conditions predicted higher alcohol use risk and longer length of opioid abstinence, respectively. In Chapter 5, smokers were classified by smaller Win-Win drift rates, while ex-smokers were classified by larger Win-Win drift-rates and smaller Lose-Lose drift-rates. In Chapter 6, RTs for Lose-Lose trials were significantly lower at T1 of nicotine abstinence compared with T0 – however drift-rate parameters did not significantly predict time to nicotine relapse.

Although the relative value between Lose-Lose and Win-Win conditions are identical, aversive versus appetitive conflict has been shown to differentially affect task performance. Cavanagh et al. (2014) postulated that slower RTs are observed on Lose-Lose conditions due to the compounding effect of two aversively conditioned stimuli, which produces slowing via the corticostriatal indirect pathway and amplifications of the STN response to conflict. Ratcliff and Frank (2012) found that the Lose-Lose condition was better fit by a DDM with delays in processing time (i.e., non-decision time) and collapsing decision boundaries, highlighting the unique effect of aversive conflict compared with other PST conditions. Consistent with these findings, the studies presented here also observed slower RTs on Lose-Lose trials during the PST test phase. However, the influence of stimulus conflict on the need to increase boundary separation was not directly tested. Despite this, these findings suggest that response to conflict in the PST test phase may show a relationship with substance-dependence. Weigard and Sripada

(2020) posited that lower trait evidence accumulation is a risk-factor for psychopathology, the current findings provide mixed support for this with regard to substance dependence (i.e., lower evidence accumulation characterises smokers, but higher rates characterise alcohol misuse and increased opioid abstinence), and suggests that relationships with drift-rate may interact with value conflict (i.e., ex-smokers are characterised by smaller drift-rates for aversive conflict, and larger drift-rates for appetitive conflict).

The relationship between conflict, boundary separation, and drift-rate has also been explored using EEG. Trial-by-trial mPFC theta during the PST test phase correlated with boundary separation in high-conflict trials in a study by Cavanagh et al. (2011), and is proposed to depend on the dorsomedial frontal cortex and STN (Frank et al., 2007). In Chapter 6, a trial-by-trial relationship was observed between theta and RT during the PST training phase, however the task design did not permit a comparison of theta effects in the test phase (i.e., the inter-trial interval between stimulus presentation and response selection was not long enough). A regression model was fit to the training phase of the PST; however separate estimates of the effect of theta on boundary separation for positive versus negative feedback did not show group differences. Further, changes in regression model parameters and EEG features across study time-points did not significantly predict time to relapse.

In Chapter 6, attenuations in FRN and fb-P3 ERPs during the training phase were observed between study time-points. The high relapse rates in Chapter 6 limited the inferences that could be made regarding this analysis (72% of participants at baseline relapsed within the first week of their smoking cessation attempt). Despite this, baseline boundary separation (from the PST test phase), baseline non-decision times (from the training and test phase), and changes in non-decision times from the training phase were selected as the most significant predictors of time to relapse in ML models. Further longitudinal research with more testing sessions scheduled within the first week of nicotine abstinence may provide insight into state-related factors predicting relapse during this critical time-period. Taken together, the results of

Chapters 5 and 6 suggest that drift-rates during the PST test phase can classify smokers, ex-smokers and non-smokers with moderate accuracy, whereas baseline boundary separation and changes in non-decision times may be related to shorter-term nicotine abstinence.

Non-decision Time

The non-decision time parameter indexes the time taken to encode stimuli and motor processes, and shows age-related slowing resulting from delays in motor execution and sensory encoding (Ratcliff et al., 2003, 2004). Smaller non-decision times have been observed during alcohol intoxication relative to sober and hangover conditions (Stock et al., 2017), in gamblers relative to controls (Wiehler & Peters, 2020), and among those on ADHD medication versus off medication compared with controls (Pederson et al., 2017). Smaller non-decision time predicted higher alcohol misuse in Chapter 3 and shorter length of opioid abstinence in Chapter 4. Larger non-decision time predicted smoker and ex-smoker group membership in Chapter 5. Non-decision times during the PST training phase were more predictive than the test phase, and showed a relationship with RTs in some cases. The high-risk alcohol group showed faster mean RTs, and the long-term opioid abstinent showed slower RTs in the PST training phase, although these did not differ significantly between groups (with the exception of the RT for EF pairs in the opioid study).

Age has been shown to affect boundary separation and non-decision time parameter estimates, with older adults displaying increased response caution and larger non-decision times compared to younger adults. Notably, age was one of the most significant predictors in ML models across all studies. However, after controlling for the effect of age, significant positive correlations were observed between boundary separation and non-decision time, therefore it is possible that shared variance between these two parameters resulted in more frequent feature selections in the ML models. The robust correlations (i.e., controlling for age) between non-decision times and length of abstinence in Chapter 4 showed a small positive correlation

distribution that was not as robust as the boundary separation parameter. In Chapter 3, the correlations between non-decision time and AUDIT score differed significantly from the null correlation distribution, showing a more robust negative relationship. It is not clear by what mechanisms non-decision times are influenced by substance-dependence/behavioural addictions, however these results, along with Wiehlers and Peters (2020), suggest a relationship with smaller non-decision times estimated from RLHDDMs.

Learning Rate

The learning rate parameters were also significant predictors in Chapters 3, 4, and 5. Learning rates represent the degree to which recent PEs affect the updating of expected rewards (i.e., the extent to which more recent or past feedback is utilised). Learning rates are sensitive to DA modulation such that higher positive learning rates have been observed in Parkinson's disease patients on versus off medication (Rutledge et al., 2009). Baker et al. (2018) found that smoking state similarly modulates learning rates during cigarette consumption versus abstinence. Nicotine abstinence has been associated with low striatal DA activity, and this was reflected in significantly lower positive learning rates and higher negative learning rates in the PST training phase (Baker et al., 2018). In Chapter 5, current smokers were classified by *reductions* in positive learning rate, relative to non-smokers. These differences are somewhat contradictory, and may be due to variations in the estimation of parameters (hierarchical combined RL and DDMs, versus Q-learning), the number of trials that models were fit to (all trials, vs. first 60 trials), and study design (cross-sectional vs. within-subjects with repeated measures).

Reduced striatal DA has also been observed in opioid-dependent individuals on methadone maintenance therapy, and mu-opioid receptor modulation has been shown to affect value-based decision-making (Eikemo et al., 2017; Liang et al., 2014). However a previous study by Myers et al. (2016) found no difference in learning rates between opioid-dependent

and healthy control groups on a reward learning task, although the former showed higher learning rates for reward that were not statistically significant. Increases in positive learning rate predicted longer length of opioid abstinence in Chapter 4. This suggests that longer length of MMT may result in functional adaptations to decision-making that improve learning from reward, however further research with a normative sample as a comparison may confirm this role of reward processing in opioid dependence.

In Chapter 3, increased negative learning rates predicted higher alcohol misuse risk. Conversely, a number of studies have shown that alcohol dependence and misuse is associated with blunted learning from negative feedback (e.g., Rossiter et al., 2012). A possible reason for inconsistencies regarding learning rate parameter predictions in the current studies is the increased variability observed for correlations between positive and negative learning and outcome variables. Repeated correlations between learning rate parameter estimates showed significantly more variability compared with the other DDM parameters (e.g., Figure 3.5), and in Chapter 3 the correlations between observed and simulated learning rates were lower than the other DDM parameters (differences between simulated and estimated parameters were also observed for higher and lower learning rates in Pedersen et al., 2017). Refining the RLHDDM to account for this increased variability may produce more reliable parameter estimates to examine individual differences in behaviour.

Self-reported Impulsivity

Impulsivity is a multidimensional construct encompassing aspects of reward seeking and disinhibition that may be considered an endophenotype for substance dependence. Self-reported impulsivity as measured by the BIS has been shown to differentiate smokers from non-smokers, with smokers showing higher attentional (Reynolds et al., 2007) and non-planning impulsivity (Flory & Manuck, 2009; Mitchell, 1999), and total BIS scores correlate positively with frequency of cigarette consumption (Skinner, Aubin, & Berlin, 2004). Here, higher non-planning

impulsivity also classified smokers relative to non-smokers, while the attentional and motor scales were non-significant predictors in Personality and Combined Models. The relationship between self-reported impulsivity among ex-smokers has not been well-characterised in the literature and studies often include small sample sizes. One study (n = 21) found lower impulsivity among ex-smokers compared with smokers, but no significant differences compared with non-smokers (Bickel, Odum, & Madden, 1999), another (n = 35) found decreased impulsivity on all BIS sub-scales compared with current smokers, however these were non-significant after controlling for age and gender, though the non-planning scale was near significant (Skinner et al., 2004). Here, ex-smokers (n = 25) were classified by decreased non-planning and increased motor impulsivity. Collectively, these findings suggest that non-planning impulsivity (i.e., the tendency to focus on the present moment with disregard for future consequences) may be relevant to successful nicotine abstinence. However, larger samples of ex-smokers are required to confirm this effect.

In Chapter 3, higher impulsivity on all BIS scales predicted higher AUDIT scores. MacKillop et al. (2016) similarly found that sub-scales of the BIS could predict low level of alcohol use risk among a large sample of young adults. A previous study with overlapping research participants from those studied in Chapter 3 has shown that non-planning and attentional impulsivity predicted alcohol intoxication frequency (O'Halloran et al., 2018). In Chapter 4, the personality models underperformed the demographic models, however lower motor impulsivity and Sensation-seeking significantly predicted longer length of abstinence in the Combined Model.

PST Summary Scores

The PST has been utilised to measure individual differences in reward versus punishment learning in a variety of clinical domains over the past 17 years. A line of research by Baker et al. (2011) observed an interaction between positive versus negative feedback learning

in the PST and substance dependence, however a replication study (Baker et al., 2013) showed that non-dependent individuals performed better overall on the PST, and no interaction effects of substance group and feedback type. In the PST Summary Models presented here, the tendency to learn from positive relative to negative feedback was indexed with Approach AC/BD variables. Here, increases in Approach AC (i.e., learning more from positive feedback in the training phase), and decreases in Approach BD (i.e., learning more from negative feedback in the training phase) predicted higher alcohol use risk and longer length of opioid dependence. In Chapters 3 and 4, increased learning from positive feedback classified high-risk alcohol users relative to low-risk, and predicted shorter length of opioid abstinence. In Chapter 5, Study 1 showed that increased learning from negative feedback classified smokers relative to non-smokers, and decreased learning from positive feedback classified smokers and ex-smokers. In Study 2, these results were consistent with the inclusion of a larger sample, and the vaper group were classified both by increased learning from positive and negative feedback. Despite the non-significant findings in Baker et al. (2013), the results from Chapter 5 show a relationship between substance dependence and positive vs. negative learning in the PST. It is worth noting however that few group differences in Approach AC/BD choices were observed between groups in the studies reported here, and that methods of analysing the PST test phase differ between studies (i.e., in Baker et al. 2011, participants were grouped into positive versus negative learning groups, which may differentially affect inferences). The computational model analyses reported here and in Baker et al. (2018) may be more useful for comparing PST-related performance across studies.

It is also worth noting that Baker et al. (2013) highlighted a lack of test-retest reliability for approach and avoidance accuracy in the PST test phase. Correlations between outcome variables measured between 7-8 weeks apart were low (r 's = .09, -.08). In Chapter 6 (Appendix D, Table D. 1), correlations between Approach AC and Approach BD at T0, T1, and T2 were considerably higher (all r 's > .4), indicating more adequate test-retest reliability within a four

week timeframe. To the best of knowledge, these are the first reported tests of test-retest reliability for the PST.

Relative utility of Computational, PST Summary, and Personality Models for predicting substance use

Across the empirical chapters, the Computational, PST Summary, and Personality Models showed a similar model fit overall. Although the test models outperformed the demographic models on average, computational features of decision-making did not predict substance group membership greater than the PST Summary or Personality model features (with the exception of Chapter 3, where the Personality models underperformed demographic models). Superior model fit was achieved by combining the computational and personality features. This is not uncommon in the literature; for example Gowin, Ball, Wittman, Tapert, and Paulus (2015) found that a personality model (including sensation-seeking, and impulsivity features) predicted methamphetamine relapse similarly to a neuroimaging model with fMRI activity recorded during a reward task. Our findings may not be surprising given a recent study by Eisenberg et al. (2019). In a data-driven approach comparing self-report and task-based factors as predictors of 'real-world' outcomes such as binge drinking, problem drinking, daily smoking, lifetime smoking, and drug-use, task-based factors were substantially worse predictors compared to self-report factors (Eisenberg et al., 2019). The task-based factors included behavioural and computational features derived from 37 experimental tasks, including the PST and DDMs. The authors suggest two possible reasons for lack of correlations between task-based dependent variables and self-regulatory behaviors: 1. these may be due to the analysis of task-based features probing cognitive functions that are irrelevant to the behavioural outcome of interest (e.g., the discounting and caution factors were the strongest task-based predictors of daily and lifetime smoking, however the 'strategic information processing' factor-including spatial span and n-back task performance DVs, did not contribute to the prediction) or 2. Discrepancies between tasks and outcomes may be as a result of theoretical overfitting;

the highly controlled nature of experimental tasks may not be ecologically valid measures of human behaviour in the real world. Of course, self-report measures are also subject to limitations regarding the individuals' knowledge about their own behaviour- alcohol and cigarette consumption frequency are estimated based on memory, and may not reflect true consumption levels. The findings of Eisenberg (2019) and others (Enkavi et al., 2019) point to the notion that latent factors derived from experimental tasks are more reliable measures of individual differences than individual dependent variables (e.g., Approach AC in the PST). Although the current study found no differences in the predictive utility of computational versus PST summary models, further research comparing the PST with a range of experimental tasks indexing temporal discounting and speeded RTs may highlight other relevant factors.

The lack of correlation between task-based (incl. computational parameters) and self-report questionnaires is relevant to research seeking to bridge various levels of analysis, as is recommended by the Research Domain Criteria (Morris & Cuthbert, 2012). A core aim of the RDoC is to link units of analysis (i.e., genes, molecules, cells, circuits, physiology, behaviour, and self-report) by researching specified constructs (e.g., reward learning) in order to better understand and treat mental illness. However, it has been shown that multiple ways of measuring the same psychological construct often do not correlate with each other, even within the same experimental task (Hedge, Bompas, & Sumner, 2020). A possible reason for lack of concordance between task-based and self-report measures is that self-report measures assess typical performance (i.e., the tendency to perform consistently across contexts), whereas task-based measures assess maximum performance (i.e., the ability to perform at a high level in one context, Wennerhold & Frieze, 2020). This may apply to forced choice decision-making tasks such as the PST that require accurate responses often within a time limit, contrasting with time-unlimited endorsement of items in questionnaires (e.g., endorsing 'I make up my mind quickly' on a Likert scale in the BIS motor impulsivity questionnaire, is likely different to the computational parameter 'impulsive information processing' indexed by RT speed in the DDM).

Computational models of task data may lead to improved reliability compared with traditional task outcomes (e.g., mean choice and RTs; Hedge et al., 2020), however beyond theoretical considerations, statistical and methodological factors have been shown to influence reliability (Brown, Chen, Gillan, & Price, 2020). For example, Dutilh et al., (2019) showed that methods of modelling reaction time data vary greatly between researchers, and that simpler models often lead to robust and accurate conclusions compared with more complex ones. Several joint computational modelling approaches to link neural and behavioural data have recently been proposed alongside methodological advances (Palestro et al., 2018), these provide useful avenues for future research linking neural covariates with decision-making behaviour.

7. 3. Methodological Strengths & Limitations

The studies described here possess a number of methodological strengths and limitations. Firstly, two versions of the PST were used in this thesis; one with a fixed number of training phase trials (Chapters 3, 5, 6), and one with training to criterion (Chapter 4). An advantage of the latter approach is that participants learn the reward/punishment contingencies to a predetermined level of accuracy, therefore improving inferences regarding response selections in the test phase (i.e., it is unlikely that stimulus selections will be random). However a challenge for research participants with cognitive impairments is that they may show consistently poor accuracy in the training phase (e.g., this has been observed among alcohol-dependent patients in Rustemeier et al., 2012). In Chapter 3, approximately half of the total former opioid dependent sample did not pass the PST training phase, which limits the generalisability of findings. An advantage of using a fixed number of trials is that participants viewed each stimulus a fixed number of times. Therefore, familiarity of stimuli is unlikely to affect selections in the test phase. Relatedly, working memory capacity and general intelligence may have affected our findings. The process of learning from stimuli and events in one's environment involves memory, as state-action values are held and updated across time. Collins and Frank (2012) suggest that working memory and reinforcement learning are two

complementary systems, with the former being constrained by capacity and sensitive to forgetting. It has been shown that RL impairments in clinical groups such as schizophrenia patients may arise due to reductions in working memory capacity rather than instrumental learning (Collins & Frank, 2014; 2017).

Secondly, state-dependent factors were not controlled for in the analysis. Emotion and mood have been demonstrated to affect learning (Eldar & Niv, 2015), while individual differences in PST performance have been observed in university samples with varying levels of self-reported anxiety and depression (Cavanagh et al., 2019), and trait anxiety has been shown to affect boundary separation in the DDM (White et al., 2010). It is conceivable that these factors may have affected participants' responding in the current studies. Craving and withdrawal may also have affected decision-making, particularly in the opioid-dependent and smoking samples. For example, Lighthall et al. (2013) found that stress modulated PST performance, such that increased stress led to diminished learning from recent feedback in the training phase and greater learning from positive feedback in the test phase. Future research may address state-dependent factors through the use of smartphone-based ecological momentary assessments, virtual reality, and wearable devices to provide relevant real-world data (Gillan & Rutledge, 2021; Johnson & Picard, 2020).

A third consideration is the difference in sample sizes between empirical chapters. Some substance-dependent samples may be challenging to recruit for laboratory-based studies (e.g., vapers and ex-smokers). Online recruitment methods may be particularly useful in targeting these samples, and have shown reliability and validity in the context of self-report measures among alcohol populations (Kim & Hodgins, 2017). Strickland and Stoops (2019) note that platforms such as Mechanical Turk can be used to complement lab-based studies in addiction science and may be suitable for longitudinal designs. Although limitations regarding random

sampling and clinical recruitment are apparent in crowdsourced studies (Mellis & Bickel, 2020), it provides an alternative to traditional methods of study recruitment.

A strength of the current analysis was the use of cross-validation procedures with out-of-sample prediction. Model performance was validated based on the prediction accuracy when applied to unseen data, which improves the generalisability of findings, and aligns with the goals of precision medicine. The use of in-sample predictions, combined with small sample sizes, presents a challenge for interpreting effect sizes in psychological research. Although internal cross-validation improves the generalisability of findings presented here, an external validation dataset would be beneficial. A second strength of the analysis was the use of an open-source code to fit HDDMs. The HDDM package (Wiecki et al., 2013) was developed with priors informed by the PST literature and has been validated in several papers to date. This contributes towards reproducibility of findings in computational modelling by allowing independent researchers to implement the same model in different environments (Poldrack et al., 2019).

7. 4. Conclusion

In conclusion, this thesis provides evidence that features of value-based decision-making are predictive of substance misuse and dependence. Across four studies, combined RL and DDMs were fit to choice and RT data from a probabilistic reward learning task, and these were used to predict various substance-dependence phenotypes. Across four studies, I sought to predict alcohol misuse risk, length of opioid abstinence, smoking group membership (i.e., current smoker, ex-smoker, vaper, non-smoker), and time to nicotine relapse during a smoking quit attempt. Four models were compared: a PST Summary model with mean choice accuracy, a Computational Model with mean parameter estimates from DDMs applied to the training and test phase, a Personality Model with self-report impulsivity, and a Combined model. The Combined Models outperformed the remaining models overall, however task-based models performed similarly to self-report models with the exception of Chapter 4. Reduced boundary separation characterised higher-risk alcohol misuse, current-smokers from non-smokers, and

shorter length of opioid abstinence, highlighting the role of response caution in predictions regarding substance dependence. Drift-rate was also a significant predictor, and differentially classified ex-smokers depending on stimulus conflict in the PST test phase. These findings highlight the utility of RLDDMs in investigating clinically relevant features of instrumental learning and decision-making, and identify features of decision-making (i.e., evidence accumulation and decision-threshold) that are relevant to substance-dependence and abstinence.

Appendix A: Chapter 3 Supplemental Results

A. 1. Correlations


Correlations between computational model parameters and the BIS and AEQ are presented below.

BIS


In the correlations between computational model parameters and the BIS, positive learning rate correlated negatively with the three 2nd order sub-scales – the remaining distributions centred on or close to 0, indicating low rho values.




RLHDDM Boundary Separation



RLHDDM Drift-rate




RLHDDM Non-decision Time






HDDM Boundary Separation








HDDM Non-decision Time






The set of correlations between the AEQ and drift-diffusion model parameters (high-risk sample N = 64) showed a positive relationship between boundary separation in the PST training phase and the arousal/power, assertiveness, positive global changes, sexual enhancement, and social/physical pleasure sub-scale. The boundary separation parameter from the test phase also showed a positive correlation with the social/physical pleasure and assertiveness sub-scales. Collectively, these suggest that increased response caution in the PST (i.e., increased boundaries or threshold to execute a decision response) is associated with increased expectancies regarding the reinforcing value of alcohol consumption among those in the high-risk alcohol group. Non-decision time in the training phase showed a positive relationship with social/physical pleasure, and drift-rate showed a positive relationship with relaxation/tension reduction. Similarly, drift-rate for high-conflict loss stimuli (LL) in the test phase showed a positive correlation with social/physical pleasure.







HDDM Boundary Separation & AEQ Correlations



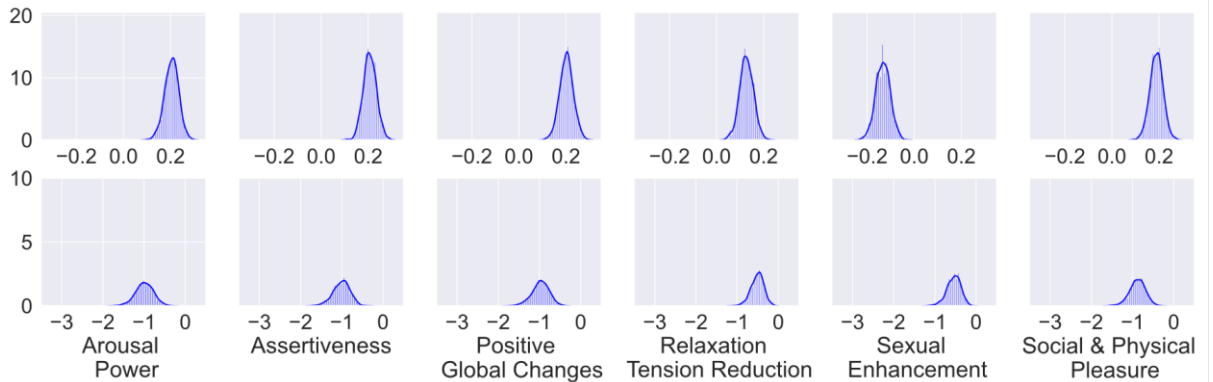
HDDM Non-decision Time & AEQ Correlations




HDDM Drift-rate (LL) & AEQ Correlations



HDDM Drift-rate (WL) & AEQ Correlations





A. 2. Non-hierarchical models

To investigate the potential effect of shrinkage towards the group posterior mean in RLHDDM parameter estimates, separate RL(H)DDM models were also fit to each subject individually.

The mean parameter estimates from these models were correlated with models hierarchically constrained by the whole group (i.e., hierarchical models were used in the subsequent machine learning models). All mean parameters from the individually estimated RLHDDM models correlated significantly ($p < .05$) with the group model (a: $\rho = 0.99$, $p < .001$; alpha neg.: $\rho = 0.68$, $p < .001$; alpha pos.: $\rho = 0.77$, $p < .001$; v: $\rho = 0.89$, $p < .001$; t: $\rho = 0.99$, $p < .001$). All mean parameters from the individually estimated HDDM models correlated significantly (all $p < .0001$) with the group model estimates (a: $\rho = 0.99$, t: $\rho = 0.99$, vLL: $\rho = 0.99$, vWL: $\rho = 0.99$, vWW: $\rho = 0.99$).

A. 3. Between-groups computational models

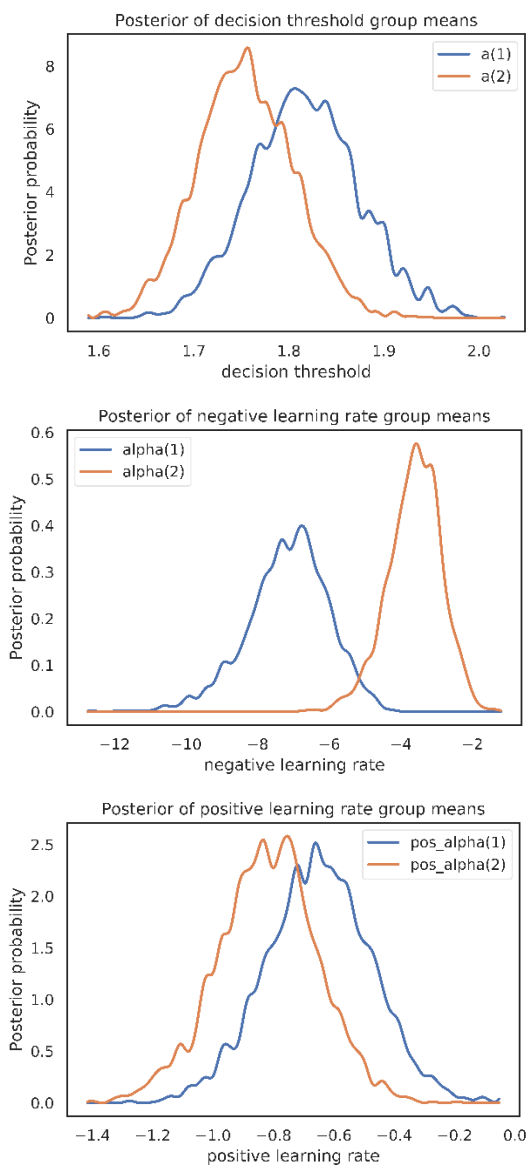
RL-HDDM

An RLHDDM model (with 15,000 samples, 3,000 burn-in, and thinning to retain every 3rd sample) was fit with separate parameters estimated for low versus high risk alcohol groups.

This model showed good convergence using the same Gelman-Rubin method (all values < 1.1 ; highest value = 1.03). The group model produced a DIC of 27141.59.

Group posteriors

Posteriors for each of the RL-HDDM parameters were compared for overlap between the high and low alcohol groups. The negative learning rate parameter for the high alcohol group was significantly higher than the low alcohol group ($q < 0.001$). The remaining group posteriors did not differ significantly between groups.







Figure 1. Posterior probability plots of parameters in the RL-HDDM compared by group, 1=low alcohol, 2=high alcohol.



Between-groups HDDM




Two types of group HDDM model were fit, three MCMCs were run for each type of model with 12,000 samples (2,000 burn-in). Model convergence was assessed using the Gelman Rubin statistic, with values <1.1 indicating convergence. In Model 1 separate parameters (a , v , t) were estimated for the low and high alcohol groups, Gelman-Rubin values for all parameters were <1.1 (max. value was 1.04), signifying adequate convergence of chains. Model fit was assessed using DIC, 21618.02. In Model 2, separate parameters were estimated for alcohol groups (a , t), and separate drift-rate parameters were additionally estimated based on the level of conflict in stimulus pairs (i.e., v_{LL} , v_{WL} , v_{WW}). LL stimulus pairs were those which included two shapes that were consistently punished in the training phase of the PST (e.g., BD, DF), WW were those pairs with two consistently rewarded shapes (e.g., AC, CE), and WL were those pairs with a combination of rewarded and punished shapes (e.g., AD, CB). Model 2

showed good convergence (max. rhat = 1.004), and better fit (DIC= 19572.36) compared with Model 1. Therefore, the group posteriors from Model 2 are presented below.

Group Posteriors

Posteriors for each of the RL-HDDM parameters were compared for overlap between the high and low alcohol groups. The parameters did not differ significantly between low (1) and high (2) alcohol groups ($q > .05$), however non-decision time, drift-rate for WL, and drift-rate for WW trended towards significance ($q = .067, .056, .072$).





A. 4. Machine Learning Models with Residual Features

In order to control for the effects of age and gender in the machine learning model, the features were adjusted using unstandardized residuals. Model 1 included the computational parameters from the RLHDDM and HDDM, and Model 2 included these parameters with the addition of the BIS scales. These were conducted with AUDIT score as the linear dependent variable (Tables 1-2).

For Model 1, the test model outperformed the demographic model in 72% of cases based on r score, and 51% of cases based on mean squared error. Mean r score for the test model (0.12 (SD=0.07)) was significantly higher for the test model compared with the demographic model

(-0.09 (0.13)), $w=6.44$, $p<001$. MSE for the test model (-32.84 (1.12)) was not significantly higher than the demographic model (-32.67 (0.97)), $p=.78$.

Table 1.

Feature	Mean Choice Frequency - Test	Mean Beta - Test	Mean Choice Frequency - demographic	Mean Beta - demographic
intercept	5	9.44348053	5	9.443
alpha_neg	4.87	0.726129206	1.57	0.0294
hddm_vWL	3.88	0.282781622	1.01	0.00257
hddm_t	3	-0.169196583	1.44	-0.0222
rlddm_a	2.34	-0.124910722	1.48	0.0203
rlddm_t	2.33	-0.119317554	1.15	0.0286
rlddm_v	1.22	0.052244757	1.48	0.0651
hddm_vWW	0.8	0.023446716	1.2	0.045
hddm_a	0.75	-0.018751271	1.2	0.05519
alpha_pos	0.69	-0.017244107	1.42	0.0380
hddm_vLL	0.66	-0.037786541	1.52	0.0686

For Model 2, with the addition of the remaining features, the test model outperformed the demographic model in 95% of cases based on r score, and 91% of cases based on mean squared error. Mean r score for the test model (0.18 (SD=0.05)) was significantly higher for the test model compared with the demographic model (-0.12 (0.12)), $w=11.76$, $p<001$. MSE for the test model (-30.85 (0.92)) was significantly higher than the demographic model (-33.25 (1.08)), $p<.001$. Mean absolute error was significantly lower for the test model (4.65 (0.06)) compared with the demographic model (4.76 (0.08)), $p<.001$.

Table 2.

Feature	Mean Choice Frequency - Test	Mean Beta - Test	Mean Choice Frequency - demographic	Mean Beta - demographic
intercept	5	9.452769047	5	9.443
BIS_2_non-planning	4.98	0.833449388	1.1	0.0219
BIS_2_motor	4.73	0.493075544	1.01	0.0089
alpha_neg	4.08	0.40046887	1.52	-0.0149
BIS_2_attentional	3.1	0.185382097	1.04	-0.0165
hddm_vWL	2.56	0.101877693	1.21	0.0013
rlddm_v	2.01	0.098001306	1.24	0.013
rlddm_a	1.96	-0.097653576	1.3	-0.0065
hddm_a	1.38	-0.059565772	1.09	0.011
hddm_vWW	1	0.032883376	1.31	0.0341
alpha_pos	0.32	0.003256638	1.29	0.0174
hddm_vLL	0.28	0.00562343	1.53	0.022

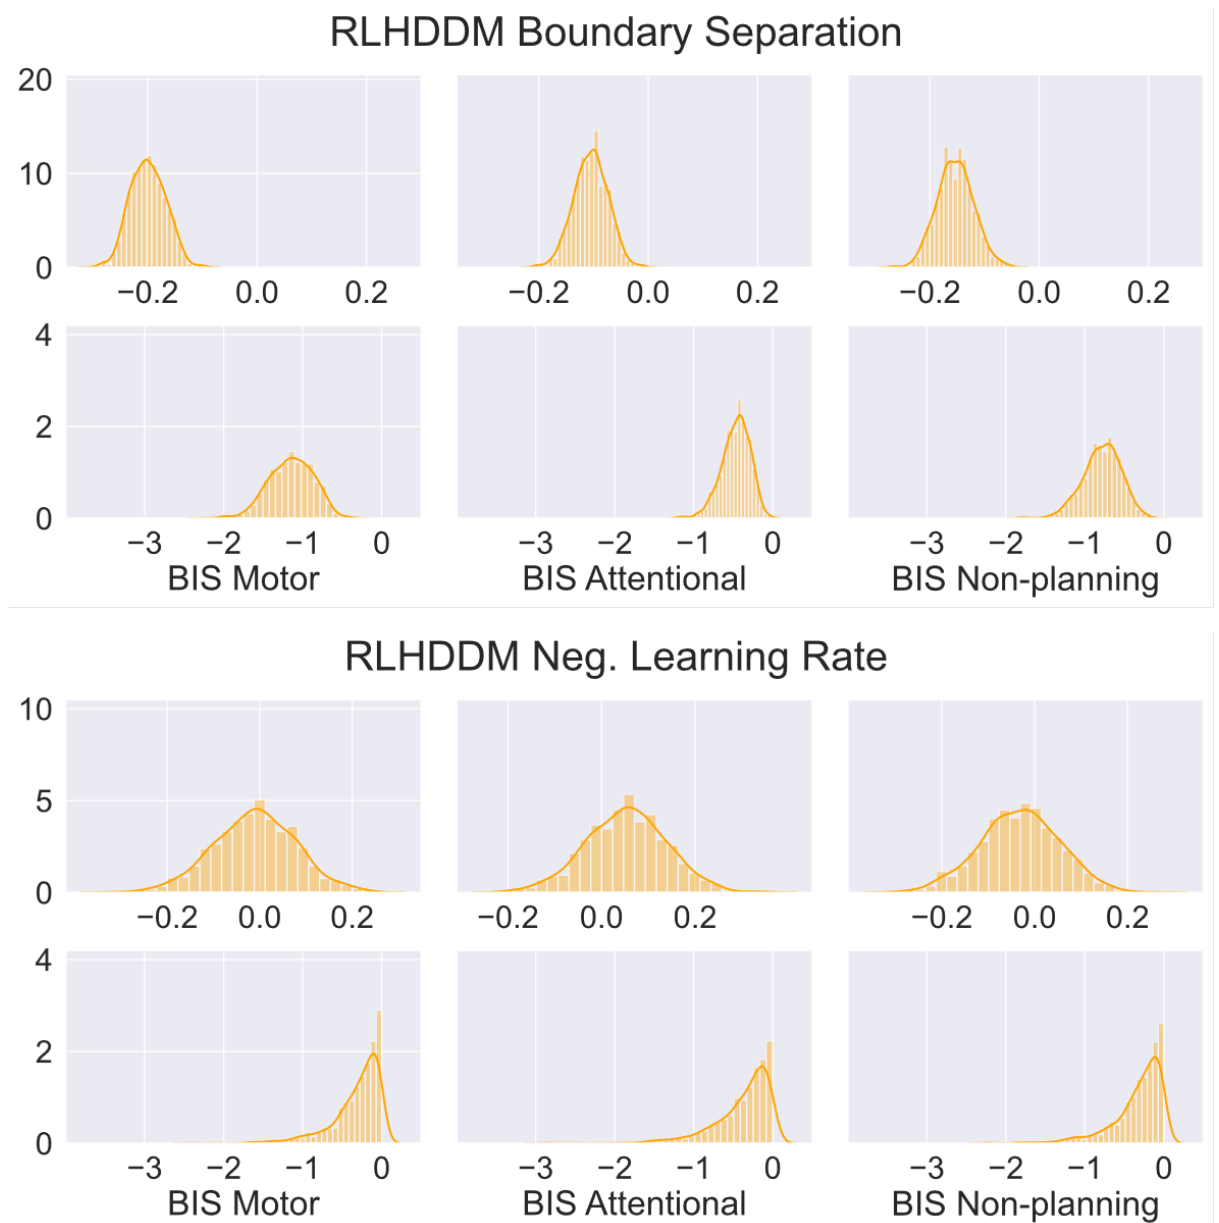
rlddm_t	0.26	-0.0060759	1.17	0.0034
hddm_t	0.15	-0.00163395	1.13	0.0174

Appendix B: Chapter 4 Supplemental Results


B. 1. Correlations

Correlations between the RLHDDM parameters and the BIS showed a negative relationship between boundary separation and the Motor and Non-planning scales, and a negative relationship between non-decision time and the Motor scale.


Figure B. 1. Distribution of rho (upper row) and $\log_{10}(\text{p-values})$ (lower row) for correlations between RLHDDM computational model parameters and the BIS.




RLHDDM Pos. Learning Rate



RLHDDM Non-decision Time







Correlations between the RLHDDM parameters and the SUPP-S showed a negative relationship between the boundary separation parameter and positive urgency, and a negative relationship between the non-decision time parameter and all SUPP-S scales. These indicate that smaller non-decision times were associated with higher positive and negative urgency, higher lack of perseverance and premeditation, and higher sensation-seeking. Narrower response thresholds (i.e., reduced caution) were associated with higher positive urgency.

Figure B. 2. Distribution of rho (upper row) and log₁₀(p-values; lower row) for correlations between RLHDDM computational model parameters and the SUPP-S.


RLHDDM Boundary Separation




RLHDDM Neg. Learning Rate




RLHDDM Pos. Learning Rate



RLHDDM Non-decision Time





B. 2. Comparison of HDDM Drift-rate Conditions

The proportion of posterior differences between the WW and LL conditions for the drift-rate parameter was significant ($q = .022$), however the differences between the remaining conditions were not considered significant ($q's > .12$). Figure B. 3. shows that drift-rates for the WW stimulus type were higher than LL and WL stimulus types.




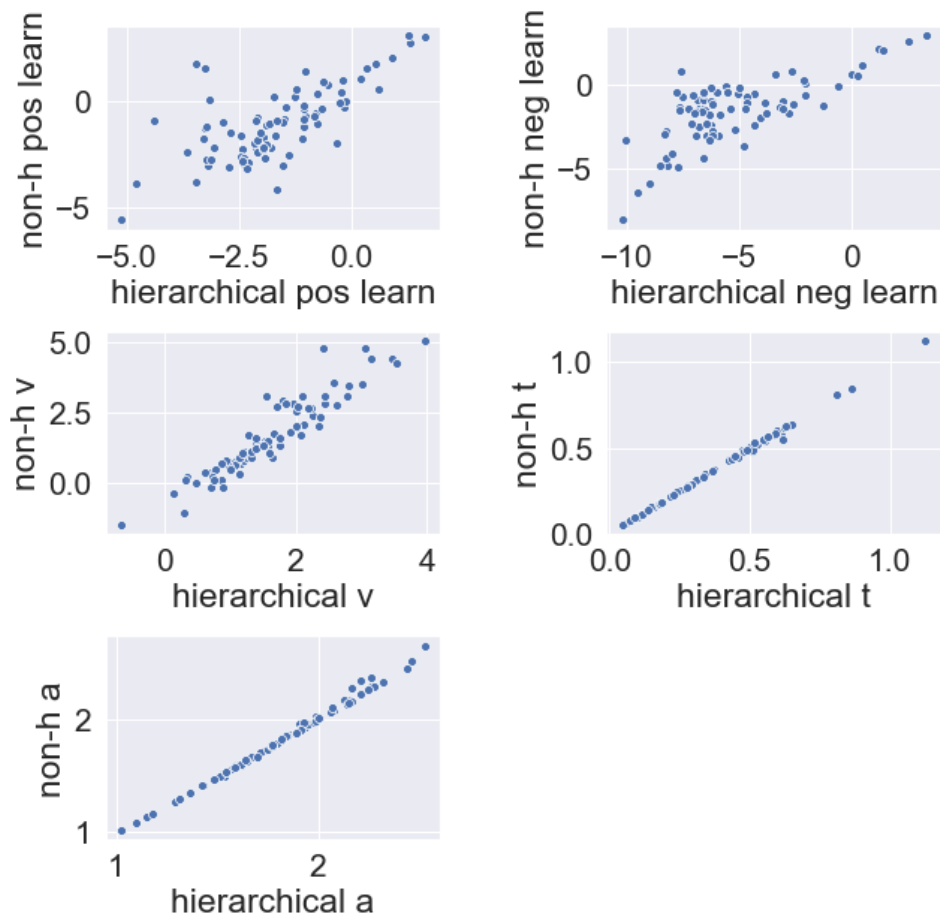
Figure B. 3. Posterior of drift-rate conflict conditions (Win-Win-, Lose-Lose, and Win-Lose) from drift-diffusion model applied to the Probabilistic Selection Task test phase.

B. 3. Non-hierarchical RLDDM

To compare the effects of potential shrinkage towards the mean in hierarchical model parameter estimates, the same analysis was conducted using RLDDM models fit to subjects

individually (i.e., non-hierarchical models). The mean parameter estimates from the hierarchical model correlated significantly with those from the non-hierarchical models (a: $\rho=0.99, p<.001$; t: $\rho=0.99, p<.001$; v: $\rho=0.95, p<.001$; negative learning rate: $\rho=0.63, p<.001$; positive learning rate: $\rho=0.66, p<.001$). Scatter plots of hierarchical vs. non-hierarchical are presented below.

Hierarchical vs. Non-hierarchical RLDDM Parameter Correlations



Appendix C: Chapter 5 Supplementary Results


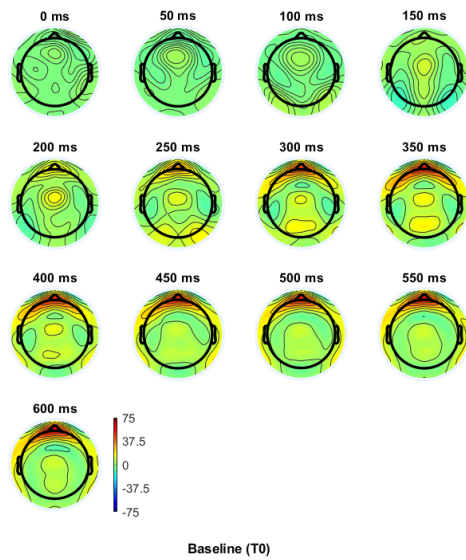
C. 1. Non-hierarchical models


These models were fit to each subject individually, to test for potential differences in estimates between hierarchical models, which are constrained by the whole group, and non-hierarchical models.

Parameter estimates from the hierarchical RLDDM and HDDM models were compared with models fit to subjects individually (i.e., non-hierarchical). Mean parameters from both types of model correlated significantly in all cases (all p 's < .001). In the RLDDM for smokers vs. non-smokers, a $\rho=0.98$, $v\rho=0.62$, $t\rho=0.99$, pos. learning rate $\rho=0.41$, and negative learning rate $\rho=0.58$. In the RLDDM model with four groups (smokers, ex-smokers, vapers, non-smokers), a $\rho=0.98$, $v\rho=0.72$, $t\rho=0.99$, pos. learning rate $\rho=0.44$, neg. learning rate $\rho = 0.61$. In the HDDM model with four groups, a $\rho=0.63$, $vLL\rho = 0.64$, $vWL\rho = 0.73$, $vWW\rho = 0.67$, and $t\rho = 0.62$.

Appendix D: Chapter 6 Supplementary Analysis


D. 1. Topoplots of EEG data






Week 4 (T2)

Figure D. 1. Topoplots of EEG activity post-feedback onset in the Probabilistic Selection Task training phase at T0, T1, and T2 follow-up sessions.



RPE+ Baseline (T0)



RPE+ Week 1 (T1)







Figure D. 2. Topoplots of EEG activity post positive feedback onset in the Probabilistic Selection Task training phase at T0, T1, and T2 follow-up sessions.



RPE- Baseline (T0)



RPE- Week 1 (T1)



RPE- Week 4 (T2)

Figure D. 3. Topoplots of EEG activity post negative feedback onset in the Probabilistic Selection Task training phase at T0, T1, and T2 follow-up sessions.

D. 1. Comparison of Computational Models

Table D. 1. 1. Summary of RLHDDMs fit to the training phase of the PST.

RLHDDM MODEL	DIC	MAX. RHAT	SAMPLES (BURN-IN)	GROUP/INDIVIDUAL
$a \sim \text{theta_next trial}$ (Fcz, 200-400 ms post-feedback)	Mean: 264.87		12,000 (2,000)	Individual
All params free to vary (a, v, t, alpha)	Mean: 219.45	1.006	12,000 (2,000)	Individual
$a \sim \text{theta:C(feedback, Treatment(0))}'$	Mean: 265.28		10,000 (2,000)	Individual


Table D. 1. 2. Summary of HDDMs fir the test phase of the PST.

HDDM MODEL	DIC	MAX. RHAT	SAMPLES (BURN-IN)	GROUP/INDIVIDUAL
All params free to vary (a, v, t)	68382.89	1.003	12,000 (2,000)	Group

<i>A, v, t depend on = study timepoint</i>	68829.72	1.007	12,000 (2,000)	Group
<i>v depends on stimulus conflict</i>	46689.15	1.003	12,000 (2,000)	Group
Regressor: A ~ study timepoint, v depends on stimulus conflict	49643.72	1.02	5,000 (500)	Group

D. 2. Within-subjects of abstinence on boundary separation

The between-subjects models on the test phase of the PST suggested that there may be differences in boundary-separation and non-decision time between study time-points. To test for within-subjects effects of study time-point (i.e., abstinence) on drift-diffusion model parameters during the PST test phase, boundary-separation was assumed to vary according to the linear model: 'a ~ C(time-point, Treatment(T0))', with baseline (T0) as the reference category.



Relative to the intercept/baseline, the regression coefficients for T1 and T2 were negative and did not overlap with zero (see below). This suggests that abstinence reduced boundary-separation on the test phase of the PST, corresponding to a reduced threshold for evidence accumulation with increased abstinence.





Figure D. 4. Within-subjects effects of changes in boundary separation across study time points. Intercept = Baseline, T1 = Week=1, T2 = Week-4.

D. 1. Supplementary Analysis of Event-related Potential-Reward Prediction Error Modulation

D. 1. 2. Reinforcement Learning

Reinforcement learning models were fit to choice data from the training phase of the PST, as described in previous research (Cavanagh et al., 2019). Based on previous findings in the literature, it was expected that an RL model with separate learning rates for positive and negative feedback would best fit the data. An RL model with a singular learning rate was formally compared with the dual learning rate model. In these models, state-action values were estimated for each stimulus type (i.e., A, B, C, D, E, F), and a softmax choice rule was used to predict the most likely action on each trial. As demonstrated in equation 1: State-action values

(Q) were updated in line with the delta learning rule with a learning rate (α) scaling the prediction error (δ).

$$Q_t = Q_{t-1} + \alpha(\delta) \quad (1)$$

Action selection probability was predicted using the softmax logistic function. Both models included a free gain adjustment parameter (β) that indicates the probability that an agent will select the most rewarding stimulus on a given trial. Lower β values indicate random choice selections, or more exploratory selections that don't conform to selections on previous trials, whereas higher values imply a greedy policy whereby the most-rewarded stimuli are consistently chosen in the task. The best-fitting model was chosen based on a combination of procedures outlined in Cavanagh et al. (2019), including Pseudo- R^2 , Akaike Information Criterion, and the likelihood ratio test to compare the relative model performance. Prediction errors from the best-fitting model were used as single trial regressors in the EEG analysis. Prediction errors were calculated as the difference between reinforcements (R) and Q-values on each trial (Q-values were initialised at 0.5 for equiprobable initial selection): $PE = R - Q$.

6.3.1.2. RL Behavioral Summary

No significant group differences in mean RL computational parameters were observed between T0 and T1 (N = 31; Table 6. 2; all p's > 0.2), or between baseline parameters for those who eventually relapsed versus those who remained abstinent (Table 6. 3; all p's > .05).

Table 6. 2. Mean reinforcement learning parameters at each study time-point.

	T0	T1	T2
POSITIVE LEARNING RATE	0.31 (0.31)	0.35 (0.31)	0.34 (0.35)
NEGATIVE LEARNING RATE	0.12 (0.23)	0.14 (0.29)	0.18 (0.29)
BETA	154.27 (429.18)	48.77 (245.39)	72.21 (306.83)

Table 6. 3. Mean reinforcement learning parameters at (i) T0 for those who eventually abstained at T1 and (ii) T0 for those who eventually relapsed at T1.

	TO_ABSTAIN	TO_RELAPSE
POSITIVE LEARNING RATE	0.32 (0.31)	0.30 (0.32)
NEGATIVE LEARNING RATE	0.09 (1.9)	0.12 (0.24)
BETA	166.42 (440.8)	149.19 (426.99)

D. 1. 3. Single trial analysis

The single trial analysis was conducted to determine if reward prediction errors significantly modulated ERPs in the time-window 0-600 ms post-feedback in the PST training phase (i.e., 308 time-points, given a sampling rate of 512Hz), similar to previously described methods (Huaser et al., 2014, Cao et al., 2021). Linear regression models identified when the RPE value predicted the ERP amplitude, and the beta weights from these models for each study participant were then subjected to a one-sample t-test. The significance threshold of the one-sample t-test was assessed by comparing t-values to null models that were constructed by performing the same analysis 1,000 times with random label permutation (i.e., shuffled RPE values within each participant). To test for group differences in RPE-ERP modulation between groups (i.e., between T0 and T1, and between baseline participants who eventually relapsed versus those who remained abstinent), t-tests with beta weights from the linear regression models were compared at each time-point (i.e 0-600ms). The significance of these t-tests were determined by constructing 1,000 null models with shuffled group labels. At each time-point, if the t-value from the true model exceeded the bottom 2.5 or top 97.5 percentile of the null models, it was deemed statistically significant.

The ERP data were transformed into current source density for the single trial analysis. Positive and negative RPE trials at the FCz electrode were entered into a multiple regression with the computationally derived trial-by-trial prediction error. The upper and lower

significance thresholds were determined using random label permutation with 1000 shuffled beta scores.

Whole Group

In the whole T0 sample, RPEs significantly modulated two ERP components (from 230 to 302ms, 37 consecutive time-points from 117 to 155; and from 324 to 466ms, 73 consecutive time-points from 166 to 239; Fig 1). At T1, two ERP components were significant, from 226 to 279ms (27 consecutive time-points from 116 to 143), and from 332 to 452ms (62 consecutive timepoints from 170 to 232; Fig 2). At T2, RPE significantly modulated ERPs from 236 to 359ms (28 consecutive time-points from 184 to 212) and from 246 to 265ms (10 consecutive time-points from 126 to 136; Fig 3).




Fig 1.

Fig. 2.





Fig. 3

Single-trial Modulation Group Comparisons

To further explore ERP/PE modulation effects, two group comparisons were made: (i) Baseline (T0) modulation for those who relapsed within the first week of their quit attempt ($n = 79$), versus those who were successfully abstinent at one-week ($n = 33$; Fig 1 & 2), and (ii) Baseline (T0) versus T1 modulation effects for abstinent participants ($N = 33$; Fig 3 & 4). Independent and paired samples t-tests were conducted to compare the beta values between groups, with 1000 shuffled group label permutations indicating the threshold of significance. To strengthen the interpretability of the effects, Bayes Factors are presented for all group comparisons (Figs. 3 & 4). No significant group differences in modulation effects were observed for both (i) and (ii).

(i)

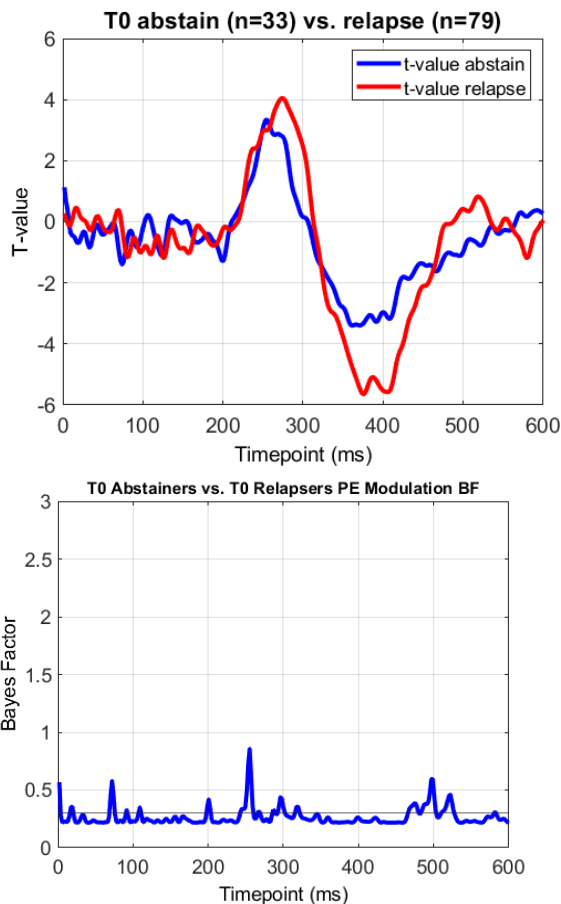


Fig. 1.





Fig. 2.

(ii)





Fig. 3.

Fig. 4.




Fig. 6

D. 3. 3. Single-trial Cox Regression

A survival analysis using Cox regression was run to predict time to relapse. Age, Gender, Marital Status, Education and Monthly Income were included as co-variates in all models. The base/null model consisted of three time-dependent variables from the MSPSS scale; Family, Friends, and Significant Other (mean C index = 0.572). A model with the time-dependent Q learning computational parameters (Positive and negative learning rates, and beta temperature) and base model variables did not outperform the base model (mean C index = 0.558). A model with the beta weights from each of the 308 timepoints of the single trial analysis did not produce a higher C index score (mean = 0.521) than the base model. A model with variables from the PSTNFB (Approach A, Approach B, Cox AC, Cox BD) performed similarly to the base model (mean C index = 0.571). This suggests that these variables did not predict time to relapse beyond the questionnaire measures.

D. 2. Correlation between PST at Study Time-points

Table D. 2. 1. Correlations between Approach AC and Approach BD at each study time-point. 0 = Baseline, 7 = Week-1, 28 = Week-4.

	Approach _AC_0	Approach _BD_0	Approach _AC_7	Approach _BD_7	Approach _AC_28	Approach _BD_28
Approach_AC_0	—					
Approach_BD_0	0.337 ***	—				
Approach_AC_7	0.471 **	0.173	—			
Approach_BD_7	0.374 *	0.500 **	0.506 **	—		
Approach_AC_28	0.580 **	0.628 **	0.913 ***	0.549 *	—	
Approach_BD_28	0.387	0.554 **	0.582 *	0.760 ***	0.503 *	—

* p < .05, ** p < .01, *** p < .001

Appendix E: Chapter 6 – Smoking Cessation Study Eligibility Criteria



Trinity College Dublin
 Coláiste na Tríonóide, Baile Átha Cliath
 The University of Dublin

Suitability Checklist

Participant ID: _____

	PLEASE CIRCLE	
Do you usually smoke at least 5-10 cigarettes per day?	YES	NO
Are you willing to quit smoking on the day after your first visit to take part in our study, and to do so for as long as possible without using any nicotine replacement products?	YES	NO
Have you ever attempted to quit smoking with professional help (not including nicotine replacement therapy)?	YES	NO
Do you own an iPhone or an Android smartphone?	YES	NO
Are you right-handed?	YES	NO
Have you ever been in an accident during which you suffered blunt force trauma to the head?	YES	NO
Have you ever been hospitalized for traumatic brain injury (this does not include concussions)?	YES	NO
Do you have Receptive Language difficulties?	YES	NO
Do you ingest/smoke cannabis regularly (twice a month)?	YES	NO
Do you currently or have you previously had an alcohol problem?	YES	NO
Do you currently or have you previously had a drug problem?	YES	NO
Do you have a learning disability? (e.g., dyslexia)	YES	NO
Have you ever been diagnosed with any general and/or specific intellectual disability?	YES	NO
Have you ever been diagnosed with a mental illness? (e.g., major depressive disorder or personality disorder)	YES	NO
Do you have a physical disability, which you feel might negatively affect your performance in this study? (e.g., motor impairment, or the effects of a stroke)	YES	NO

Inclusion criteria

- Current daily smoker: self-reported as having smoked at least 100 cigarettes in their lifetime, and currently smoking at least 5 cigarettes a day for at least the last 12 months (without any quit attempts during the past week)
- Baseline Carbon Monoxide (CO) \geq 5 parts per million (ppm) and a cotinine score \geq 1 (NicAlert reading)
- Between 18 and 70 years old
- English fluency (native speakers or a score $>$ 16 on an online test: <https://www.cambridgeenglish.org/test-your-english/general-english/>)
- Willing to come into Trinity College Dublin to take part in three laboratory sessions
- Normal or corrected to normal vision (self-reported)
- Normal or corrected to normal hearing (self-reported)
- Willing to quit smoking in the next 30 days
- An iPhone (iOS8 or above) or an Android (4.0.3 or above) smartphone and have mobile internet access

Exclusion criteria

- For women, (self-reported) currently pregnant or lactating
- Consuming any recreational drugs besides tobacco/alcohol and cannabis (less than twice a week): self-reported
- Currently or have previously had problems controlling alcohol/drug consumption: self-reported
- Currently taking prescription drugs other than SSRI's: self-reported
- Learning or language disabilities (e.g., dyslexia or receptive language difficulties): self-reported
- Physical disability (e.g., motor impairment, or the effects of a stroke)
- Any previous experience of traumatic head injury and/ or loss of consciousness: self-reported
- Self-reported history of neurological (e.g., MS, Parkinson's or previous stroke) or psychological diagnosis

References

- Aberg, K. C., Doell, K. C., & Schwartz, S. (2016). Linking Individual Learning Styles to Approach Avoidance Motivational Traits and Computational Aspects of Reinforcement Learning. *PloS One*, 11(11), e0166675. <https://doi.org/10.1371/journal.pone.0166675>
- Acion, L., Kelmansky, D., van der Laan, M., Sahker, E., Jones, D., & Arndt, S. (2017). Use of a machine learning framework to predict substance use disorder treatment success. *PloS One*, 12(4), e0175383. <https://doi.org/10.1371/journal.pone.0175383>
- Adams, C. D., & Dickinson, A. (1981). Instrumental Responding following Reinforcer Devaluation. *The Quarterly Journal of Experimental Psychology Section B*, 33(2b), 109–121. <https://doi.org/10.1080/14640748108400816>
- Ahmed, S. H. (2008). The origin of addictions by means of unnatural decision. *Behavioral and Brain Sciences*, 31(4), 437–438. <https://doi.org/10.1017/S0140525X08004731>
- Ahn, W.-Y., & Vassileva, J. (2016). Machine-learning identifies substance-specific behavioral markers for opiate and stimulant dependence. *Drug and Alcohol Dependence*, 161, 247–257. <https://doi.org/10.1016/j.drugalcdep.2016.02.008>
- Ahn, W.-Y., Vasilev, G., Lee, S.-H., Busemeyer, J. R., Kruschke, J. K., Bechara, A., & Vassileva, J. (2014). Decision-making in stimulant and opiate addicts in protracted abstinence: Evidence from computational modeling with pure users. *Frontiers in Psychology*, 5. <https://doi.org/10.3389/fpsyg.2014.00849>
- American Psychiatric Association (Ed.). (2013). *Diagnostic and statistical manual of mental disorders* (5th ed.). Washington, DC: Author.
- Baker, T. B., Piper, M. E., McCarthy, D. E., Majeskie, M. R., & Fiore, M. C. (2004). Addiction motivation reformulated: an affective processing model of negative reinforcement. *Psychological Review*, 111(1), 33–51. <https://doi.org/10.1037/0033-295X.111.1.33>

- Baker, T. E., Stockwell, T., & Holroyd, C. B. (2013). Constraints on decision making: Implications from genetics, personality, and addiction. *Cognitive, Affective, & Behavioral Neuroscience*, 13(3), 417–436. <https://doi.org/10.3758/s13415-013-0164-8>
- Baker, T. E., Stockwell, T., Barnes, G., & Holroyd, C. B. (2011). Individual differences in substance dependence: at the intersection of brain, behaviour and cognition: Substance dependence. *Addiction Biology*, 16(3), 458–466. <https://doi.org/10.1111/j.1369-1600.2010.00243.x>
- Baker, T. E., Zeighami, Y., Dagher, A., & Holroyd, C. B. (2018). Smoking decisions: Altered reinforcement learning signals induced by nicotine state. *Nicotine & Tobacco Research: Official Journal of the Society for Research on Nicotine and Tobacco*. <https://doi.org/10.1093/ntr/nty136>
- Barbeau, A. M., Burda, J., & Siegel, M. (2013). Perceived efficacy of e-cigarettes versus nicotine replacement therapy among successful e-cigarette users: A qualitative approach. *Addiction Science & Clinical Practice*, 8(1), 5. <https://doi.org/10.1186/1940-0640-8-5>
- Barr, R. S., Pizzagalli, D. A., Culhane, M. A., Goff, D. C., & Evins, A. E. (2008). A Single Dose of Nicotine Enhances Reward Responsiveness in Nonsmokers: Implications for Development of Dependence. *Biological Psychiatry*, 63(11), 1061–1065. <https://doi.org/10.1016/j.biopsych.2007.09.015>
- Bauer, L. O. (2001). Predicting relapse to alcohol and drug abuse via quantitative electroencephalography. *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology*, 25(3), 332–340. [https://doi.org/10.1016/S0893-133X\(01\)00236-6](https://doi.org/10.1016/S0893-133X(01)00236-6)
- Beaton, L. E., Azma, S., & Marinkovic, K. (2018). When the brain changes its mind: Oscillatory dynamics of conflict processing and response switching in a flanker task during alcohol challenge. *PLoS One*, 13(1), e0191200. <https://doi.org/10.1371/journal.pone.0191200>
- Bechara, A., Berridge, K. C., Bickel, W. K., Morón, J. A., Williams, S. B., & Stein, J. S. (2019). A Neurobehavioral Approach to Addiction: Implications for the Opioid Epidemic and the

Psychology of Addiction. *Psychological Science in the Public Interest*, 20(2), 96–127.

<https://doi.org/10.1177/1529100619860513>

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1-3), 7-15. doi: 10.1016/0010-0277(94)90018-3

Bentzley, B. S., Barth, K. S., Back, S. E., & Book, S. W. (2015). Discontinuation of buprenorphine maintenance therapy: Perspectives and outcomes. *Journal of Substance Abuse Treatment*, 52, 48–57. <https://doi.org/10.1016/j.jsat.2014.12.011>

Benureau, F. C. Y., & Rougier, N. P. (2018). Re-run, Repeat, Reproduce, Reuse, Replicate: Transforming Code into Scientific Contributions. *Frontiers in Neuroinformatics*, 11. <https://doi.org/10.3389/fninf.2017.00069>

Beylergil, S. B., Beck, A., Deserno, L., Lorenz, R. C., Rapp, M. A., Schlagenhaut, F., Heinz, A., & Obermayer, K. (2017). Dorsolateral prefrontal cortex contributes to the impaired behavioral adaptation in alcohol dependence. *NeuroImage : Clinical*, 15, 80–94. <https://doi.org/10.1016/j.nicl.2017.04.010>

Białaszek, W., Marcowski, P., & Cox, D. J. (2017). Differences in Delay, but not Probability Discounting, in Current Smokers, E-cigarette Users, and Never Smokers. *Psychological Record*, 67(2), 223–230. Scopus. <https://doi.org/10.1007/s40732-017-0244-1>

Bickel, W. K., & Marsch, L. A. (2001). Toward a behavioral economic understanding of drug dependence: Delay discounting processes. *Addiction (Abingdon, England)*, 96(1), 73–86. <https://doi.org/10.1046/j.1360-0443.2001.961736.x>

Bickel, W. K., Odum, A. L., & Madden, G. J. (1999). Impulsivity and cigarette smoking: Delay discounting in current, never, and ex-smokers. *Psychopharmacology*, 146(4), 447–454. <https://doi.org/10.1007/pl00005490>

Bickel, Warren K., Mellis, A. M., Snider, S. E., Athamneh, L. N., Stein, J. S., & Pope, D. A. (2018). 21st century neurobehavioral theories of decision making in addiction: Review and evaluation.

Pharmacology, Biochemistry, and Behavior, 164, 4–21.

<https://doi.org/10.1016/j.pbb.2017.09.009>

Biernacki, K., McLennan, S. N., Terrett, G., Labuschagne, I., & Rendell, P. G. (2016). Decision-making ability in current and past users of opiates: A meta-analysis. *Neuroscience and Biobehavioral Reviews*, 71, 342–351. <https://doi.org/10.1016/j.neubiorev.2016.09.011>

Blum, K., Braverman, E. R., Holder, J. M., Lubar, J. F., Monastera, V. J., Miller, D., Lubar, J. O., Chen, T. J., Comings, D. E. (2000). Reward deficiency syndrome: a biogenetic model for the diagnosis and treatment of impulsive, addictive, and compulsive behaviors. *J Psychoactive Drugs* 32:1–112.

Bódi, N., Kéri, S., Nagy, H., Moustafa, A., Myers, C. E., Daw, N., ... Gluck, M. A. (2009). Reward-learning and the novelty-seeking personality: a between- and within-subjects study of the effects of dopamine agonists on young Parkinson's patients. *Brain: A Journal of Neurology*, 132(Pt 9), 2385–2395. <https://doi.org/10.1093/brain/awp094>

Boehm, U., Marsman, M., Matzke, D., & Wagenmakers, E.-J. (2018). On the importance of avoiding shortcuts in applying cognitive models to hierarchical data. *Behavior Research Methods*, 50(4), 1614–1631. <https://doi.org/10.3758/s13428-018-1054-3>

Botvinik-Nezer, R., Holzmeister, F., Camerer, C. F., Dreber, A., Huber, J., Johannesson, M., Kirchler, M., Iwanir, R., Mumford, J. A., Adcock, R. A., Avesani, P., Baczkowski, B. M., Bajracharya, A., Bakst, L., Ball, S., Barilari, M., Bault, N., Beaton, D., Beitner, J., ... Schonberg, T. (2020). Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582(7810), 84–88. <https://doi.org/10.1038/s41586-020-2314-9>

Brand, M., Roth-Bauer, M., Driessen, M., & Markowitsch, H. J. (2008). Executive functions and risky decision-making in patients with opiate dependence. *Drug and Alcohol Dependence*, 97(1), 64–72. <https://doi.org/10.1016/j.drugalcdep.2008.03.017>

Bridwell, D. A., Cavanagh, J. F., Collins, A. G. E., Nunez, M. D., Srinivasan, R., Stober, S., & Calhoun, V. D. (2018). Moving Beyond ERP Components: A Selective Review of Approaches to

Integrate EEG and Behavior. *Frontiers in Human Neuroscience*, 12.

<https://doi.org/10.3389/fnhum.2018.00106>

Briggs, Z., O'Connor, M., Jollans, E. K., O'Halloran, L., Dymond, S., & Whelan, R. (2015). Flexible emotion-based decision-making behavior varies in current and former smokers. *Addictive Behaviors*, 45, 269–275. <https://doi.org/10.1016/j.addbeh.2015.02.011>

Brody, A. L., Olmstead, R. E., London, E. D., Farahi, J., Meyer, J. H., Grossman, P., Lee, G. S., Huang, J., Hahn, E. L., & Mandelkern, M. A. (2004). Smoking-Induced Ventral Striatum Dopamine Release. *American Journal of Psychiatry*, 161(7), 1211–1218.

<https://doi.org/10.1176/appi.ajp.161.7.1211>

Brown, V. M., Chen, J., Gillan, C. M., & Price, R. B. (2020). Improving the Reliability of Computational Analyses: Model-Based Planning and Its Relationship With Compulsivity. *Biological Psychiatry. Cognitive Neuroscience and Neuroimaging*, 5(6), 601–609.

<https://doi.org/10.1016/j.bpsc.2019.12.019>

Browning, M., Carter, C. S., Chatham, C., Ouden, H. D., Gillan, C. M., Baker, J. T., Chekroud, A. M., Cools, R., Dayan, P., Gold, J., Goldstein, R. Z., Hartley, C. A., Kepecs, A., Lawson, R. P., Mourao-Miranda, J., Phillips, M. L., Pizzagalli, D. A., Powers, A., Rindskopf, D., ... Paulus, M. (2020). Realizing the Clinical Potential of Computational Psychiatry: Report From the Banbury Center Meeting, February 2019. *Biological Psychiatry*, 88(2), e5–e10.

<https://doi.org/10.1016/j.biopsych.2019.12.026>

Bu, J., Ma, R., Fan, C., Sun, S., Cheng, Y., Piao, Y., Zhang, P., Liu, C., & Zhang, X. (2019). Low-Theta Electroencephalography Coherence Predicts Cigarette Craving in Nicotine Addiction. *Frontiers in Psychiatry*, 10. <https://doi.org/10.3389/fpsy.2019.00296>

Büchel, C., Peters, J., Banaschewski, T., Bokde, A. L. W., Bromberg, U., Conrod, P. J., Flor, H., Papadopoulos, D., Garavan, H., Gowland, P., Heinz, A., Walter, H., Ittermann, B., Mann, K., Martinot, J.-L., Paillère-Martinot, M.-L., Nees, F., Paus, T., Pausova, Z., ... Knutson, B. (2017). Blunted ventral striatal responses to anticipated rewards foreshadow problematic drug use

in novelty-seeking adolescents. *Nature Communications*, 8(1), 14140.

<https://doi.org/10.1038/ncomms14140>

Bühler, M., Vollstädt-Klein, S., Kobiella, A., Budde, H., Reed, L. J., Braus, D. F., Büchel, C., & Smolka, M. N. (2010). Nicotine dependence is characterized by disordered reward processing in a network driving motivation. *Biological Psychiatry*, 67(8), 745–752.

<https://doi.org/10.1016/j.biopsych.2009.10.029>

Butler, K., Rusted, J., Gard, P., & Jackson, A. (2017). Performance monitoring in nicotine dependence: Considering integration of recent reinforcement history. *Pharmacology, Biochemistry, and Behavior*, 156, 63–70. <https://doi.org/10.1016/j.pbb.2017.04.004>

Byrne, K. A., & Worthy, D. A. (2019). Examining the link between reward and response inhibition in individuals with substance abuse tendencies. *Drug and Alcohol Dependence*, 194, 518–525.

Scopus. <https://doi.org/10.1016/j.drugalcdep.2018.11.014>

Bzdok, D., & Meyer-Lindenberg, A. (2018). Machine Learning for Precision Psychiatry: Opportunities and Challenges. *Biological Psychiatry. Cognitive Neuroscience and Neuroimaging*, 3(3), 223–230. <https://doi.org/10.1016/j.bpsc.2017.11.007>

Cahill, K., Stevens, S., Perera, R., & Lancaster, T. (2013). Pharmacological interventions for smoking cessation: An overview and network meta-analysis. *Cochrane Database of Systematic Reviews*, 5. <https://doi.org/10.1002/14651858.CD009329.pub2>

Cahill, K., Stevens, S., Perera, R., & Lancaster, T. (2013). Pharmacological interventions for smoking cessation: An overview and network meta-analysis. *Cochrane Database of Systematic Reviews*, 5. <https://doi.org/10.1002/14651858.CD009329.pub2>

Cao, Z., Bennett, M., O'Halloran, L., Pragulbickaite, G., Flanagan, L., McHugh, L., & Whelan, R. (2021). Aberrant reward prediction errors in young adult at-risk alcohol users. *Addiction Biology*, 26(1), e12873. <https://doi.org/10.1111/adb.12873>

Cao, Z., Ottino-Gonzalez, J., Cupertino, R. B., Schwab, N., Hoke, C., Catherine, O., Cousijn, J., Dagher, A., Foxe, J. J., Goudriaan, A. E., Hester, R., Hutchison, K., Li, C.-S. R., London, E. D.,

- Lorenzetti, V., Luijten, M., Martin-Santos, R., Momenan, R., Paulus, M. P., ... Garavan, H. (n.d.). Mapping cortical and subcortical asymmetries in substance dependence: Findings from the ENIGMA Addiction Working Group. *Addiction Biology*, *n/a(n/a)*, e13010. <https://doi.org/10.1111/adb.13010>
- Carballo, M., & López, V. (2014). Error monitoring recovery after abstinence and treatment in cocaine addicts / Recuperación del procesamiento de errores tras abstinencia y tratamiento en adictos a cocaína. *Estudios de Psicología*, *35*(1), 193–199. <https://doi.org/10.1080/02109395.2014.893656>
- Caswell, A. J., Celio, M. A., Morgan, M. J., & Duka, T. (2015). Impulsivity as a multifaceted construct related to excessive drinking among UK students. *Alcohol and alcoholism*, *51*(1), 77-83.
- Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences*, *18*(8), 414–421. <https://doi.org/10.1016/j.tics.2014.04.012>
- Cavanagh, J. F., & Shackman, A. J. (2015). Frontal Midline Theta Reflects Anxiety and Cognitive Control: Meta-Analytic Evidence. *Journal of Physiology, Paris*, *109*(0), 3–15. <https://doi.org/10.1016/j.jphysparis.2014.04.003>
- Cavanagh, J. F., Bismark, A. W., Frank, M. J., & Allen, J. J. B. (2019). Multiple Dissociations Between Comorbid Depression and Anxiety on Reward and Punishment Processing: Evidence From Computationally Informed EEG. *Computational Psychiatry (Cambridge, Mass.)*, *3*, 1–17. https://doi.org/10.1162/cpsy_a_00024
- Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., & Frank, M. J. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience*, *14*(11), 1462–1467. <https://doi.org/10.1038/nn.2925>
- Cavanagh, J. F., Wiecki, T. V., Kochar, A., & Frank, M. J. (2014). Eye tracking and pupillometry are indicators of dissociable latent decision processes. *Journal of Experimental Psychology. General*, *143*(4), 1476–1488. <https://doi.org/10.1037/a0035813>

- Chase, H. W., Frank, M. J., Michael, A., Bullmore, E. T., Sahakian, B. J., & Robbins, T. W. (2010). Approach and avoidance learning in patients with major depression and healthy controls: relation to anhedonia. *Psychological Medicine*, 40(3), 433–440. <https://doi.org/10.1017/S0033291709990468>
- Chen, H., Mojtahedzadeh, N., Belanger, M. J., Nebe, S., Kuitunen-Paul, S., Sebold, M., Garbusow, M., Huys, Q. J. M., Heinz, A., Rapp, M. A., & Smolka, M. N. (2021). Model-based and model-free control predicts alcohol consumption developmental trajectory in young adults – a three-year prospective study. *Biological Psychiatry*. <https://doi.org/10.1016/j.biopsych.2021.01.009>
- Chivers, L. L., Hand, D. J., Priest, J. S., & Higgins, S. T. (2016). E-cigarette use among women of reproductive age: Impulsivity, cigarette smoking status, and other risk factors. *Preventive Medicine*, 92, 126–134. <https://doi.org/10.1016/j.ypmed.2016.07.029>
- Chowdhury, R., Guitart-Masip, M., Lambert, C., Dayan, P., Huys, Q., Düzel, E., & Dolan, R. J. (2013). Dopamine restores reward prediction errors in old age. *Nature Neuroscience*, 16(5), 648–653. <https://doi.org/10.1038/nn.3364>
- Cicero, D. C., Martin, E. A., Becker, T. M., & Kerns, J. G. (2014). Reinforcement learning deficits in people with schizophrenia persist after extended trials. *Psychiatry Research*, 220(3), 760–764. <https://doi.org/10.1016/j.psychres.2014.08.013>
- Cicero, D. C., Martin, E. A., Becker, T. M., & Kerns, J. G. (2014). Reinforcement learning deficits in people with schizophrenia persist after extended trials. *Psychiatry Research*, 220(3), 760–764. <https://doi.org/10.1016/j.psychres.2014.08.013>
- Clapp, W. C., & Gazzaley, A. (2012). Distinct mechanisms for the impact of distraction and interruption on working memory in aging. *Neurobiology of Aging*, 33(1), 134–148. <https://doi.org/10.1016/j.neurobiolaging.2010.01.012>
- Cohen, M. X. (2017). Where Does EEG Come From and What Does It Mean? *Trends in Neurosciences*, 40(4), 208–218. <https://doi.org/10.1016/j.tins.2017.02.004>

- Collins, A. G. E., & Frank, M. J. (2012). How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *The European Journal of Neuroscience*, *35*(7), 1024–1035. <https://doi.org/10.1111/j.1460-9568.2011.07980.x>
- Collins, A. G. E., Albrecht, M. A., Waltz, J. A., Gold, J. M., & Frank, M. J. (2017). Interactions Among Working Memory, Reinforcement Learning, and Effort in Value-Based Choice: A New Paradigm and Selective Deficits in Schizophrenia. *Biological Psychiatry*, *82*(6), 431–439. <https://doi.org/10.1016/j.biopsych.2017.05.017>
- Collins, A. G. E., Brown, J. K., Gold, J. M., Waltz, J. A., & Frank, M. J. (2014). Working Memory Contributions to Reinforcement Learning Impairments in Schizophrenia. *The Journal of Neuroscience*, *34*(41), 13747–13756. <https://doi.org/10.1523/JNEUROSCI.0989-14.2014>
- Connor, J. P., Symons, M., Feeney, G. F. X., Young, R. M., & Wiles, J. (2007). The application of machine learning techniques as an adjunct to clinical decision making in alcohol dependence treatment. *Substance Use & Misuse*, *42*(14), 2193–2206. <https://doi.org/10.1080/10826080701658125>
- Corr, P. J. (2004). Reinforcement sensitivity theory and personality. *Neuroscience & Biobehavioral Reviews*, *28*(3), 317–332. <https://doi.org/10.1016/j.neubiorev.2004.01.005>
- Coughlin, L. N., Tegge, A. N., Sheffer, C. E., & Bickel, W. K. (2018). A Machine-Learning Approach to Predicting Smoking Cessation Treatment Outcomes. *Nicotine & Tobacco Research*, *22*(3), 415–422. <https://doi.org/10.1093/ntr/nty259>
- Dani, J. A., & Harris, R. A. (2005). Nicotine addiction and comorbidity with alcohol abuse and mental illness. *Nature Neuroscience*, *8*(11), 1465–1470. <https://doi.org/10.1038/nn1580>
- Darke, S., Sims, J., McDonald, S., & Wickes, W. (2000). Cognitive impairment among methadone maintenance patients. *Addiction (Abingdon, England)*, *95*(5), 687–695. <https://doi.org/10.1046/j.1360-0443.2000.9556874.x>

- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*(6), 1204–1215.
<https://doi.org/10.1016/j.neuron.2011.02.027>
- Dayan, P., & Berridge, K. C. (2014). Model-based and model-free Pavlovian reward learning: Revaluation, revision, and revelation. *Cognitive, Affective & Behavioral Neuroscience*, *14*(2), 473–492. <https://doi.org/10.3758/s13415-014-0277-8>
- de Ruiter, M. B., Veltman, D. J., Goudriaan, A. E., Oosterlaan, J., Sjoerds, Z., & van den Brink, W. (2009). Response Perseveration and Ventral Prefrontal Sensitivity to Reward and Punishment in Male Problem Gamblers and Smokers. *Neuropsychopharmacology*, *34*(4), 1027–1038. <https://doi.org/10.1038/npp.2008.175>
- Decker, J. H., Otto, A. R., Daw, N. D., & Hartley, C. A. (2016). From Creatures of Habit to Goal-Directed Learners: Tracking the Developmental Emergence of Model-Based Reinforcement Learning. *Psychological Science*, *27*(6), 848–858.
<https://doi.org/10.1177/0956797616639301>
- Dick, D. M., Smith, G., Olausson, P., Mitchell, S. H., Leeman, R. F., O'Malley, S. S., & Sher, K. (2010). Understanding the construct of impulsivity and its relationship to alcohol use disorders. *Addiction Biology*, *15*(2), 217–226. <https://doi.org/10.1111/j.1369-1600.2009.00190.x>
- Dinga, R., Schmaal, L., Penninx, B. W. J. H., Veltman, D. J., & Marquand, A. F. (2020). Controlling for effects of confounding variables on machine learning predictions. *BioRxiv*, 2020.08.17.255034.
<https://doi.org/10.1101/2020.08.17.255034>
- Donaldson, K. R., Ait Oumeziane, B., Hélie, S., & Foti, D. (2016). The temporal dynamics of reversal learning: P3 amplitude predicts valence-specific behavioral adjustment. *Physiology & Behavior*, *161*, 24–32. <https://doi.org/10.1016/j.physbeh.2016.03.034>
- Doran, N., & Tully, L. (2018). Impulsivity and tobacco product use over time. *Addictive Behaviors*, *85*, 153–157. Scopus. <https://doi.org/10.1016/j.addbeh.2018.06.007>

- Dowd, E. C., Frank, M. J., Collins, A., Gold, J. M., & Barch, D. M. (2016). Probabilistic Reinforcement Learning in Patients With Schizophrenia: Relationships to Anhedonia and Avolition. *Biological Psychiatry. Cognitive Neuroscience and Neuroimaging*, 1(5), 460–473. <https://doi.org/10.1016/j.bpsc.2016.05.005>
- Dully, J., McGovern, D. P., & O’Connell, R. G. (2018). The impact of natural aging on computational and neural indices of perceptual decision making: A review. *Behavioural Brain Research*, 355, 48–55. <https://doi.org/10.1016/j.bbr.2018.02.001>
- Dutilh, G., Annis, J., Brown, S. D., Cassey, P., Evans, N. J., Grasman, R. P. P. P., Hawkins, G. E., Heathcote, A., Holmes, W. R., Kryptos, A.-M., Kupitz, C. N., Leite, F. P., Lerche, V., Lin, Y.-S., Logan, G. D., Palmeri, T. J., Starns, J. J., Trueblood, J. S., van Maanen, L., ... Donkin, C. (2019). The Quality of Response Time Data Inference: A Blinded, Collaborative Assessment of the Validity of Cognitive Models. *Psychonomic Bulletin & Review*, 26(4), 1051–1069. <https://doi.org/10.3758/s13423-017-1417-2>
- Dwyer, D. B., Falkai, P., & Koutsouleris, N. (2018). Machine Learning Approaches for Clinical Psychology and Psychiatry. *Annual Review of Clinical Psychology*, 14(1), 91–118. <https://doi.org/10.1146/annurev-clinpsy-032816-045037>
- Eaton, N. R., Rodriguez-Seijas, C., Carragher, N., & Krueger, R. F. (2015). Transdiagnostic factors of psychopathology and substance use disorders: A review. *Social Psychiatry and Psychiatric Epidemiology*, 50(2), 171–182. Scopus. <https://doi.org/10.1007/s00127-014-1001-2>
- Eikemo, M., Biele, G., Willoch, F., Thomsen, L., & Leknes, S. (2017). Opioid Modulation of Value-Based Decision-Making in Healthy Humans. *Neuropsychopharmacology*, 42(9), 1833–1840. <https://doi.org/10.1038/npp.2017.58>
- Eisenberg, I. W., Bissett, P. G., Zeynep Enkavi, A., Li, J., MacKinnon, D. P., Marsch, L. A., & Poldrack, R. A. (2019). Uncovering the structure of self-regulation through data-driven ontology discovery. *Nature Communications*, 10(1), 2319. <https://doi.org/10.1038/s41467-019-10301-1>

- Ekhtiari, H., Victor, T. A., & Paulus, M. P. (2017). Aberrant decision-making and drug addiction—How strong is the evidence? *Current Opinion in Behavioral Sciences*, *13*, 25–33.
<https://doi.org/10.1016/j.cobeha.2016.09.002>
- Eldar, E., & Niv, Y. (2015). Interaction between emotional state and learning underlies mood instability. *Nature Communications*, *6*(1), 6149. <https://doi.org/10.1038/ncomms7149>
- Endrass, T., Kloft, L., Kaufmann, C., & Kathmann, N. (2011). Approach and avoidance learning in obsessive-compulsive disorder. *Depression and Anxiety*, *28*(2), 166–172.
<https://doi.org/10.1002/da.20772>
- Eppinger, B., & Kray, J. (2011). To choose or to avoid: age differences in learning from positive and negative feedback. *Journal of Cognitive Neuroscience*, *23*(1), 41–52.
<https://doi.org/10.1162/jocn.2009.21364>
- Eppinger, B., Kray, J., Mock, B., & Mecklinger, A. (2008). Better or worse than expected? Aging, learning, and the ERN. *Neuropsychologia*, *46*(2), 521–539.
<https://doi.org/10.1016/j.neuropsychologia.2007.09.001>
- Ersche, K. D., Meng, C., Ziauddeen, H., Stochl, J., Williams, G. B., Bullmore, E. T., & Robbins, T. W. (2020). Brain networks underlying vulnerability and resilience to drug addiction. *Proceedings of the National Academy of Sciences*, *117*(26), 15253–15261.
<https://doi.org/10.1073/pnas.2002509117>
- Ersche, K. D., Meng, C., Ziauddeen, H., Stochl, J., Williams, G. B., Bullmore, E. T., & Robbins, T. W. (2020). Brain networks underlying vulnerability and resilience to drug addiction. *Proceedings of the National Academy of Sciences*, *117*(26), 15253–15261.
<https://doi.org/10.1073/pnas.2002509117>
- Ersche, K. D., Williams, G. B., Robbins, T. W., & Bullmore, E. T. (2013). Meta-analysis of structural brain abnormalities associated with stimulant drug dependence and neuroimaging of addiction vulnerability and resilience. *Current Opinion in Neurobiology*, *23*(4), 615–624.
<https://doi.org/10.1016/j.conb.2013.02.017>

- Eshel, N., & Roiser, J. P. (2010). Reward and punishment processing in depression. *Biological Psychiatry*, 68(2), 118–124. <https://doi.org/10.1016/j.biopsych.2010.01.027>
- European Monitoring Centre for Drugs and Addiction, (2019). European Drug Report 2019: Trends and Developments. Available from:
https://www.emcdda.europa.eu/publications/edr/trends-developments/2019_en
- Evans, D. E., Maxfield, N. D., Van Rensburg, K. J., Oliver, J. A., Jentink, K. G., & Drobles, D. J. (2013). Nicotine deprivation influences P300 markers of cognitive control. *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology*, 38(12), 2525–2531. <https://doi.org/10.1038/npp.2013.159>
- Evans, D. E., Sutton, S. K., Oliver, J. A., & Drobles, D. J. (2015). Cortical Activity Differs During Nicotine Deprivation versus Satiation in Heavy Smokers. *Psychopharmacology*, 232(11), 1879–1885. <https://doi.org/10.1007/s00213-014-3821-x>
- Everitt B. J., Dickinson A, Robbins TW. (2001). The neuropsychological basis of addictive behaviour. *Brain Research. Brain Research Reviews*. 36(2–3):129–138.
- Everitt B. J., Robbins T. W. (2005). Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nature Neuroscience*. 8(11):1481–1489.
- Evren, C., Alniak, I., Karabulut, V., Cetin, T., Umut, G., Agachanli, R., & Evren, B. (2018). Relationship of probable ADHD with novelty seeking, severity of psychopathology and borderline personality disorder in a sample of patients with opioid use disorder. *Psychiatry and Clinical Psychopharmacology*, 28(1), 48–55. Scopus.
<https://doi.org/10.1080/24750573.2017.1395312>
- Fehr, C., Yakushev, I., Hohmann, N., Buchholz, H.-G., Landvogt, C., Deckers, H., ... Schreckenberger, M. (2008). Association of low striatal dopamine d2 receptor availability with nicotine dependence similar to that seen with other drugs of abuse. *The American Journal of Psychiatry*, 165(4), 507–514. <https://doi.org/10.1176/appi.ajp.2007.07020352>

- Finkel, D., Reynolds, C. A., McArdle, J. J., & Pedersen, N. L. (2007). Age changes in processing speed as a leading indicator of cognitive aging. *Psychology and Aging, 22*(3), 558–568.
<https://doi.org/10.1037/0882-7974.22.3.558>
- Flory, J. D., & Manuck, S. B. (2009). Impulsiveness and cigarette smoking. *Psychosomatic Medicine, 71*(4), 431–437. <https://doi.org/10.1097/PSY.0b013e3181988c2d>
- Fontanesi, L., Gluth, S., Spektor, M. S., & Rieskamp, J. (2019). A reinforcement learning diffusion decision model for value-based decisions. *Psychonomic Bulletin & Review, 26*(4), 1099–1121. <https://doi.org/10.3758/s13423-018-1554-2>
- Forstmann, B. U., Tittgemeyer, M., Wagenmakers, E.-J., Derrfuss, J., Imperati, D., & Brown, S. (2011). The Speed-Accuracy Tradeoff in the Elderly Brain: A Structural Model-Based Approach. *Journal of Neuroscience, 31*(47), 17242–17249.
<https://doi.org/10.1523/JNEUROSCI.0309-11.2011>
- Frank, M. J., & Hutchison, K. (2009). Genetic contributions to avoidance-based decisions: striatal D2 receptor polymorphisms. *Neuroscience, 164*(1), 131–140.
<https://doi.org/10.1016/j.neuroscience.2009.04.048>
- Frank, M. J., & Kong, L. (2008). Learning to avoid in older age. *Psychology and Aging, 23*(2), 392–398. <https://doi.org/10.1037/0882-7974.23.2.392>
- Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. (2015). fMRI and EEG Predictors of Dynamic Decision Parameters during Human Reinforcement Learning. *The Journal of Neuroscience, 35*(2), 485.
<https://doi.org/10.1523/JNEUROSCI.2036-14.2015>
- Frank, M. J., Moustafa, A. A., Haughey, H. M., Curran, T., & Hutchison, K. E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proceedings of the National Academy of Sciences, 104*(41), 16311–16316.
<https://doi.org/10.1073/pnas.0706111104>

- Frank, M. J., Santamaria, A., O'Reilly, R. C., & Willcutt, E. (2007). Testing computational models of dopamine and noradrenaline dysfunction in attention deficit/hyperactivity disorder. *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology*, 32(7), 1583–1599. <https://doi.org/10.1038/sj.npp.1301278>
- Frank, M. J., Seeberger, L. C., & O'reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in parkinsonism. *Science (New York, N.Y.)*, 306(5703), 1940–1943. <https://doi.org/10.1126/science.1102941>
- Frässle, S., Marquand, A. F., Schmaal, L., Dinga, R., Veltman, D. J., van der Wee, N. J. A., van Tol, M.-J., Schöbi, D., Penninx, B. W. J. H., & Stephan, K. E. (2020). Predicting individual clinical trajectories of depression with generative embedding. *NeuroImage: Clinical*, 26, 102213. <https://doi.org/10.1016/j.nicl.2020.102213>
- Friston, K. J., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: The brain as a phantastic organ. *The Lancet. Psychiatry*, 1(2), 148–158. [https://doi.org/10.1016/S2215-0366\(14\)70275-5](https://doi.org/10.1016/S2215-0366(14)70275-5)
- Galandra, C., Basso, G., Cappa, S., & Canessa, N. (2018). The alcoholic brain: Neural bases of impaired reward-based decision-making in alcohol use disorders. *Neurological Sciences*, 39(3), 423–435. <https://doi.org/10.1007/s10072-017-3205-1>
- Garavan, H., Brennan, K. L., Hester, R., & Whelan, R. (2013). The Neurobiology of Successful Abstinence. *Current Opinion in Neurobiology*, 23(4), 668–674. <https://doi.org/10.1016/j.conb.2013.01.029>
- Garbusow, M., Schad, D. J., Sebold, M., Friedel, E., Bernhardt, N., Koch, S. P., Steinacher, B., Kathmann, N., Geurts, D. E. M., Sommer, C., Müller, D. K., Nebe, S., Paul, S., Wittchen, H.-U., Zimmermann, U. S., Walter, H., Smolka, M. N., Sterzer, P., Rapp, M. A., ... Heinz, A. (2016). Pavlovian-to-instrumental transfer effects in the nucleus accumbens relate to relapse in alcohol dependence. *Addiction Biology*, 21(3), 719–731. <https://doi.org/10.1111/adb.12243>

- Garland, E. L., & Howard, M. O. (2014). Opioid attentional bias and cue-elicited craving predict future risk of prescription opioid misuse among chronic pain patients. *Drug and Alcohol Dependence*, *144*, 283–287. <https://doi.org/10.1016/j.drugalcdep.2014.09.014>
- GBD 2015 Tobacco Collaborators. (2017). Smoking prevalence and attributable disease burden in 195 countries and territories, 1990-2015: A systematic analysis from the Global Burden of Disease Study 2015. *Lancet (London, England)*, *389*(10082), 1885–1906. [https://doi.org/10.1016/S0140-6736\(17\)30819-X](https://doi.org/10.1016/S0140-6736(17)30819-X)
- Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *ELife*, *5*, e11305. <https://doi.org/10.7554/eLife.11305>
- Gillan, C. M., & Whelan, R. (2017). What big data can do for treatment in psychiatry. *Current Opinion in Behavioral Sciences*, *18*, 34–42. <https://doi.org/10.1016/j.cobeha.2017.07.003>
- Ginsburg, G. S., & Phillips, K. A. (2018). Precision Medicine: From Science to Value. *Health Affairs (Project Hope)*, *37*(5), 694–701. <https://doi.org/10.1377/hlthaff.2017.1624>
- Goudie, A. J., Field, M., & Cole, J. (2008). Different vulnerabilities for addiction may contribute to the same phenomena and some additional interactions. *Behavioral and Brain Sciences*, *31*(4), 445–446. <https://doi.org/10.1017/S0140525X08004810>
- Gowin, J. L., Ball, T. M., Wittmann, M., Tapert, S. F., & Paulus, M. P. (2015). Individualized relapse prediction: Personality measures and striatal and insular activity during reward-processing robustly predict relapse. *Drug and Alcohol Dependence*, *152*, 93–101. <https://doi.org/10.1016/j.drugalcdep.2015.04.018>
- Gradin, V. B., Baldacchino, A., Balfour, D., Matthews, K., & Steele, J. D. (2014). Abnormal brain activity during a reward and loss task in opiate-dependent patients receiving methadone maintenance therapy. *Neuropsychopharmacology*, *39*(4), 885–894. Scopus. <https://doi.org/10.1038/npp.2013.289>

- Grant, J. E., Lust, K., Fridberg, D. J., King, A. C., & Chamberlain, S. R. (2019). E-Cigarette Use (Vaping) is Associated with Illicit Drug Use, Mental Health Problems, and Impulsivity in University Students. *Annals of Clinical Psychiatry : Official Journal of the American Academy of Clinical Psychiatrists*, *31*(1), 27–35.
- Griswold, M. G., Fullman, N., Hawley, C., Arian, N., Zimsen, S. R. M., Tymeson, H. D., Venkateswaran, V., Tapp, A. D., Forouzanfar, M. H., Salama, J. S., Abate, K. H., Abate, D., Abay, S. M., Abbafati, C., Abdulkader, R. S., Abebe, Z., Aboyans, V., Abrar, M. M., Acharya, P., ... Gakidou, E. (2018). Alcohol use and burden for 195 countries and territories, 1990–2016: A systematic analysis for the Global Burden of Disease Study 2016. *The Lancet*, *392*(10152), 1015–1035. [https://doi.org/10.1016/S0140-6736\(18\)31310-2](https://doi.org/10.1016/S0140-6736(18)31310-2)
- Grogan, J. P., Tsivos, D., Smith, L., Knight, B. E., Bogacz, R., Whone, A., & Coulthard, E. J. (2017). Effects of dopamine on reinforcement learning and consolidation in Parkinson’s disease. *ELife*, *6*. <https://doi.org/10.7554/eLife.26801>
- Gueguen, M. C., Schweitzer, E. M., & Konova, A. B. (2021). Computational theory-driven studies of reinforcement learning and decision-making in addiction: What have I learned? *Current Opinion in Behavioral Sciences*, *38*, 40–48. <https://doi.org/10.1016/j.cobeha.2020.08.007>
- Harlé, K. M., Stewart, J. L., Zhang, S., Tapert, S. F., Yu, A. J., & Paulus, M. P. (2015). Bayesian neural adjustment of inhibitory control predicts emergence of problem stimulant use. *Brain*, *138*(11), 3413–3426. <https://doi.org/10.1093/brain/awv246>
- Harper, J., Malone, S. M., & Iacono, W. G. (2018a). Conflict-related medial frontal theta as an endophenotype for alcohol use disorder. *Biological Psychology*, *139*, 25–38. <https://doi.org/10.1016/j.biopsycho.2018.10.002>
- Harper, J., Malone, S. M., & Iacono, W. G. (2018b). Impact of alcohol use on EEG dynamics of response inhibition: A cotwin control analysis. *Addiction Biology*, *23*(1), 256–267. <https://doi.org/10.1111/adb.12481>

- Hartz, S. M., & Bierut, L. J. (2010). Genetics of addictions. *Clinics in Laboratory Medicine*, 30(4), 847–864. <https://doi.org/10.1016/j.cll.2010.07.005>
- Hedge, C., Bompas, A., & Sumner, P. (2020). Task Reliability Considerations in Computational Psychiatry. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 5(9), 837–839. <https://doi.org/10.1016/j.bpsc.2020.05.004>
- Hedge, C., Powell, G., Bompas, A., & Sumner, P. (2020). Self-reported impulsivity does not predict response caution. *Personality and Individual Differences*, 167, 110257. <https://doi.org/10.1016/j.paid.2020.110257>
- Holcomb, L. A., Huang, S., Cruz, S. M., & Marinkovic, K. (2019). Neural oscillatory dynamics of inhibitory control in young adult binge drinkers. *Biological Psychology*, 146, 107732. <https://doi.org/10.1016/j.biopsycho.2019.107732>
- Houston, R. J., & Schliez, N. J. (2018). Event-Related Potentials as Biomarkers of Behavior Change Mechanisms in Substance Use Disorder Treatment. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 3(1), 30–40. <https://doi.org/10.1016/j.bpsc.2017.09.006>
- Hsieh, L.-T., & Ranganath, C. (2014). Frontal midline theta oscillations during working memory maintenance and episodic encoding and retrieval. *NeuroImage*, 85, 721–729. Scopus. <https://doi.org/10.1016/j.neuroimage.2013.08.003>
- Huettel, S. A., Stowe, C. J., Gordon, E. M., Warner, B. T., & Platt, M. L. (2006). Neural Signatures of Economic Preferences for Risk and Ambiguity. *Neuron*, 49(5), 765–775. <https://doi.org/10.1016/j.neuron.2006.01.024>
- Hughes, J. R., Budney, A. J., Muellers, S. R., Lee, D. C., Callas, P. W., Sigmon, S. C., Fingar, J. R., & Priest, J. (2017). Does Tobacco Abstinence Decrease Reward Sensitivity? A Human Laboratory Test. *Nicotine & Tobacco Research*, 19(6), 677–685. <https://doi.org/10.1093/ntr/ntw204>
- Huhn, A., Meyer, R., Harris, J., Ayaz, H., Deneke, E., Stankoski, D., & Bunce, S. (2016). Evidence of Anhedonia and Differential Reward Processing in Prefrontal Cortex Among Post-

- Withdrawal Patients with Prescription Opiate Dependence. *Brain Research Bulletin*, 123, 102–109. <https://doi.org/10.1016/j.brainresbull.2015.12.004>
- Hull, C. L. (1943). *Principles of behavior: An introduction to behavior theory*. New York: D. Appleton-Century Company, Incorporated.
- Huys, Q. J. M., Browning, M., Paulus, M. P., & Frank, M. J. (2020). Advances in the computational understanding of mental illness. *Neuropsychopharmacology*, 1–17. <https://doi.org/10.1038/s41386-020-0746-4>
- Huys, Q. J. M., Maia, T. V., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience*, 19(3), 404–413. <https://doi.org/10.1038/nn.4238>
- Inchley, J., World Health Organization, & Regional Office for Europe. (2018). *Adolescent alcohol-related behaviours: Trends and inequalities in the WHO European Region, 2002-2014*. World Health Organization, Regional Office for Europe.
- Inquisit 5 [Computer software]. (2016). Retrieved from <https://www.millisecond.com>.
- Insel T, Cuthbert B, Garvey M, Heinssen R, Pine DS, Quinn K, Wang P. (2010). Research domain criteria (RDoC): toward a new classification framework for research on mental disorders. *The American Journal of Psychiatry*. 167(7):748–751.
- James, S. L., Abate, D., Abate, K. H., Abay, S. M., Abbafati, C., Abbasi, N., Abbastabar, H., Abd-Allah, F., Abdela, J., Abdelalim, A., Abdollahpour, I., Abdulkader, R. S., Abebe, Z., Abera, S. F., Abil, O. Z., Abraha, H. N., Abu-Raddad, L. J., Abu-Rmeileh, N. M. E., Accrombessi, M. M. K., ... Murray, C. J. L. (2018). Global, regional, and national incidence, prevalence, and years lived with disability for 354 diseases and injuries for 195 countries and territories, 1990–2017: A systematic analysis for the Global Burden of Disease Study 2017. *The Lancet*, 392(10159), 1789–1858. [https://doi.org/10.1016/S0140-6736\(18\)32279-7](https://doi.org/10.1016/S0140-6736(18)32279-7)

- Jean-Richard-dit-Bressel, P., Ma, C., Bradfield, L. A., Killcross, S., & McNally, G. P. (2019). Punishment insensitivity emerges from impaired contingency detection, not aversion insensitivity or reward dominance. *ELife*, *8*, e52765. <https://doi.org/10.7554/eLife.52765>
- Kale, D., Pickering, A., & Cooper, A. (2020). Examining the relationship between impulsivity-related personality traits and e-cigarette use in adults. *Addictive Behaviors*, *106*, 106348. <https://doi.org/10.1016/j.addbeh.2020.106348>
- Kamarajan, C., Pandey, A. K., Chorlian, D. B., Manz, N., Stimus, A. T., Anokhin, A. P., Bauer, L. O., Kuperman, S., Kramer, J., Bucholz, K. K., Schuckit, M. A., Hesselbrock, V. M., & Porjesz, B. (2015). Deficient Event-Related Theta Oscillations in Individuals at Risk for Alcoholism: A Study of Reward Processing and Impulsivity Features. *PLoS ONE*, *10*(11). <https://doi.org/10.1371/journal.pone.0142659>
- Kanen, J. W., Ersche, K. D., Fineberg, N. A., Robbins, T. W., & Cardinal, R. N. (2019). Computational modelling reveals contrasting effects on reinforcement learning and cognitive flexibility in stimulant use disorder and obsessive-compulsive disorder: Remediating effects of dopaminergic D2/3 receptor agents. *Psychopharmacology*, *236*(8), 2337–2358. <https://doi.org/10.1007/s00213-019-05325-w>
- Kendler KS (2005). Toward a philosophical structure for psychiatry. *American Journal of Psychiatry* *162*(3), 433–40.
- Kendler KS (2008). Explanatory models for psychiatric illness. *American Journal of Psychiatry* *165*(6), 695–702.
- Kiiski, H., Jollans, L., Donnchadha, S. Ó., Nolan, H., Lonergan, R., Kelly, S., O'Brien, M. C., Kinsella, K., Bramham, J., Burke, T., Hutchinson, M., Tubridy, N., Reilly, R. B., & Whelan, R. (2018). Machine Learning EEG to Predict Cognitive Functioning and Processing Speed Over a 2-Year Period in Multiple Sclerosis Patients and Controls. *Brain Topography*, *31*(3), 346–363. <https://doi.org/10.1007/s10548-018-0620-4>

- Kim, H. S., & Hodgins, D. C. (2017). Reliability and validity of data obtained from alcohol, cannabis, and gambling populations on Amazon's Mechanical Turk. *Psychology of Addictive Behaviors: Journal of the Society of Psychologists in Addictive Behaviors*, 31(1), 85–94.
<https://doi.org/10.1037/adb0000219>
- Konova, A. B., Lopez-Guzman, S., Urmanche, A., Ross, S., Louie, K., Rotrosen, J., & Glimcher, P. W. (2020). Computational Markers of Risky Decision-making for Identification of Temporal Windows of Vulnerability to Opioid Use in a Real-world Clinical Setting. *JAMA Psychiatry*, 77(4), 368. <https://doi.org/10.1001/jamapsychiatry.2019.4013>
- Koob, G. F. (2009). Neurobiological substrates for the dark side of compulsivity in addiction. *Neuropharmacology*, 56 Suppl 1, 18–31.
<https://doi.org/10.1016/j.neuropharm.2008.07.043>
- Koob, G. F. (2020). Neurobiology of Opioid Addiction: Opponent Process, Hyperkatifeia, and Negative Reinforcement. *Biological Psychiatry*, 87(1), 44–53.
<https://doi.org/10.1016/j.biopsych.2019.05.023>
- Koob, G. F., & Volkow, N. D. (2010). Neurocircuitry of addiction. *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology*, 35(1), 217–238.
<https://doi.org/10.1038/npp.2009.110>
- Kriegler, J., Wegener, S., Richter, F., Scherbaum, N., Brand, M., & Wegmann, E. (2019). Decision making of individuals with heroin addiction receiving opioid maintenance treatment compared to early abstinent users. *Drug and Alcohol Dependence*, 205. Scopus.
<https://doi.org/10.1016/j.drugalcdep.2019.107593>
- Kühn, S., Schubert, F., & Gallinat, J. (2010). Reduced thickness of medial orbitofrontal cortex in smokers. *Biological Psychiatry*, 68(11), 1061–1065.
<https://doi.org/10.1016/j.biopsych.2010.08.004>
- Kunisato, Y., Okamoto, Y., Ueda, K., Onoda, K., Okada, G., Yoshimura, S., ... Yamawaki, S. (2012). Effects of depression on reward-based decision making and variability of action in

probabilistic learning. *Journal of Behavior Therapy and Experimental Psychiatry*, 43(4), 1088–1094. <https://doi.org/10.1016/j.jbtep.2012.05.007>

Lane, S. D., Cherek, D. R., Rhodes, H. M., Pietras, C. J., & Tcheremissine, O. V. (2003). Relationships among laboratory and psychometric measures of impulsivity: Implications in substance abuse and dependence. *Addictive Disorders & Their Treatment*, 2(2), 33–40. <https://doi.org/10.1097/00132576-200302020-00001>

Lannoy, S., Billieux, J., Dormal, V., & Maurage, P. (n.d.). Behavioral and Cerebral Impairments Associated with Binge Drinking in Youth: A Critical Review. *Psychologica Belgica*, 59(1), 116–155. <https://doi.org/10.5334/pb.476>

Lanza, H. I., Motlagh, G., & Orozco, M. (2020). E-cigarette use among young adults: A latent class analysis examining co-use and correlates of nicotine vaping. *Addictive Behaviors*, 110, 106528. <https://doi.org/10.1016/j.addbeh.2020.106528>

Lawson, R. P., Mathys, C., & Rees, G. (2017). Adults with autism overestimate the volatility of the sensory environment. *Nature Neuroscience*, 20(9), 1293–1299. <https://doi.org/10.1038/nn.4615>

Le Strat, Y., Ramoz, N., & Gorwood, P. (2010). In Alcohol-Dependent Drinkers, What Does the Presence of Nicotine Dependence Tell Us About Psychiatric and Addictive Disorders Comorbidity? *Alcohol and Alcoholism*, 45(2), 167–172. <https://doi.org/10.1093/alcalc/agg094>

Lee, B. D. (2018). Ten simple rules for documenting scientific software. *PLOS Computational Biology*, 14(12), e1006561. <https://doi.org/10.1371/journal.pcbi.1006561>

Lee, Y. H., Bang, H., & Kim, D. J. (2016). How to Establish Clinical Prediction Models. *Endocrinology and Metabolism (Seoul, Korea)*, 31(1), 38–44. <https://doi.org/10.3803/EnM.2016.31.1.38>

Lespine, L. F., Rueda-Delgado, L. M., Vahey, N., Kiiski, H., Enz, N., Boyle, R., ... Whelan, R. (2020). Predictors Of Successful Nicotine Abstinence Using A Machine Learning Approach:

Genetics, Environment And Neurocognitive Endophenotypes. Retrieved from:

<https://osf.io/dympt/>

Lespine, L. F., Rueda-Delgado, L. M., Vahey, N., Kiiski, H., Enz, N., Boyle, R., ... Whelan, R. (Under Review). Changes in inhibition-related brain function and psychological flexibility during smoking abstinence predict longer time to relapse.

LI, H., ZHOU, Y., LI, S., WANG, Q., PAN, L., YANG, X., ... JIA, C. (2015). The Relationship between Nicotine Dependence and Age among Current Smokers. *Iranian Journal of Public Health*, 44(4), 495–500.

Li, J., Weidacker, K., Mandali, A., Zhang, Y., Whiteford, S., Ren, Q., Zhou, Z., Zhou, H., Jiang, H., Du, J., Zhang, C., Sun, B., & Voon, V. (2021). Impulsivity and craving in subjects with opioid use disorder on methadone maintenance treatment. *Drug and Alcohol Dependence*, 219, 108483. <https://doi.org/10.1016/j.drugalcdep.2020.108483>

Li, Q., Li, W., Wang, H., Wang, Y., Zhang, Y., Zhu, J., Zheng, Y., Zhang, D., Wang, L., Li, Y., Yan, X., Chang, H., Fan, M., Li, Z., Tian, J., Gold, M. S., Wang, W., & Liu, Y. (2015). Predicting subsequent relapse by drug-related cue-induced brain activation in heroin addiction: An event-related functional magnetic resonance imaging study. *Addiction Biology*, 20(5), 968–978. <https://doi.org/10.1111/adb.12182>

Lighthall, N. R., Pearson, J. M., Huettel, S. A., & Cabeza, R. (2018). Feedback-based learning in aging: Contributions and trajectories of change in striatal and hippocampal systems. *Journal of Neuroscience*, 0769–18. <https://doi.org/10.1523/JNEUROSCI.0769-18.2018>

Limongi, R., Bohaterewicz, B., Nowicka, M., Plewka, A., & Friston, K. J. (2018). Knowing when to stop: Aberrant precision and evidence accumulation in schizophrenia. *Schizophrenia Research*, 197, 386–391. <https://doi.org/10.1016/j.schres.2017.12.018>

Liverant, G. I., Sloan, D. M., Pizzagalli, D. A., Harte, C. B., Kamholz, B. W., Rosebrock, L. E., Cohen, A. L., Fava, M., & Kaplan, G. B. (2014). Associations among smoking, anhedonia, and reward

learning in depression. *Behavior Therapy*, 45(5), 651–663.

<https://doi.org/10.1016/j.beth.2014.02.004>

Loganathan, K., & Ho, E. T. W. (2021). Value, drug addiction and the brain. *Addictive Behaviors*, 116,

106816. <https://doi.org/10.1016/j.addbeh.2021.106816>

Lopez-Quintero, C., Hasin, D. S., de Los Cobos, J. P., Pines, A., Wang, S., Grant, B. F., & Blanco, C.

(2011). Probability and predictors of remission from life-time nicotine, alcohol, cannabis or cocaine dependence: Results from the National Epidemiologic Survey on Alcohol and

Related Conditions. *Addiction (Abingdon, England)*, 106(3), 657–669.

<https://doi.org/10.1111/j.1360-0443.2010.03194.x>

Low, E. C. T., Ong, M. C. C., & Tan, M. (2004). Breath carbon monoxide as an indication of smoking habit in the military setting. *Singapore Medical Journal*, 45, 578–582.

Luijten, M., Gillan, C. M., de Wit, S., Franken, I. H. A., Robbins, T. W., & Ersche, K. D. (2020). Goal-Directed and Habitual Control in Smokers. *Nicotine & Tobacco Research*, 22(2), 188–195.

<https://doi.org/10.1093/ntr/ntz001>

Ly, A., Boehm, U., Heathcote, A., Turner, B. M., Forstmann, B., Marsman, M., & Matzke, D. (2017). A

Flexible and Efficient Hierarchical Bayesian Approach to the Exploration of Individual

Differences in Cognitive-model-based Neuroscience. In A. A. Moustafa (Ed.), *Computational Models of Brain and Behavior* (pp. 467–479). John Wiley & Sons, Ltd.

<https://doi.org/10.1002/9781119159193.ch34>

Mackey, S., Allgaier, N., Charani, B., Spechler, P., Orr, C., Bunn, J., Allen, N. B., Alia-Klein, N.,

Batalla, A., Blaine, S., Brooks, S., Caparelli, E., Chye, Y. Y., Cousijn, J., Dagher, A., Desrivieres,

S., Feldstein-Ewing, S., Foxe, J. J., Goldstein, R. Z., ... ENIGMA Addiction Working Group.

(2019). Mega-Analysis of Gray Matter Volume in Substance Dependence: General and

Substance-Specific Regional Effects. *The American Journal of Psychiatry*, 176(2), 119–128.

<https://doi.org/10.1176/appi.ajp.2018.17040415>

- MacKillop, J., Weafer, J., C Gray, J., Oshri, A., Palmer, A., & de Wit, H. (2016). The latent structure of impulsivity: Impulsive choice, impulsive action, and impulsive personality traits. *Psychopharmacology*, 233(18), 3361–3370. <https://doi.org/10.1007/s00213-016-4372-0>
- Madden, G. J., Petry, N. M., Badger, G. J., & Bickel, W. K. (1997). Impulsive and self-control choices in opioid-dependent patients and non-drug-using control participants: Drug and monetary rewards. *Experimental and Clinical Psychopharmacology*, 5(3), 256–262. <https://doi.org/10.1037//1064-1297.5.3.256>
- Maia, T. V., & Frank, M. J. (2011). From reinforcement learning models to psychiatric and neurological disorders. *Nature Neuroscience*, 14(2), 154–162. <https://doi.org/10.1038/nn.2723>
- Mak, K. K., Lee, K., & Park, C. (2019). Applications of machine learning in addiction studies: A systematic review. *Psychiatry Research*, 275, 53–60. <https://doi.org/10.1016/j.psychres.2019.03.001>
- Mallard, T. T., Ashenurst, J. R., Harden, K. P., & Fromme, K. (2018). GABRA2, alcohol, and illicit drug use: An event-level model of genetic risk for polysubstance use. *Journal of Abnormal Psychology*, 127(2), 190–201. <https://doi.org/10.1037/abn0000333>
- Mandali, A., Weidacker, K., Kim, S.-G., & Voon, V. (2019). The ease and sureness of a decision: Evidence accumulation of conflict and uncertainty. *Brain*, 142(5), 1471–1482. <https://doi.org/10.1093/brain/awz013>
- Maril, S., Hassin-Baer, S., Cohen, O. S., & Tomer, R. (2013). Effects of asymmetric dopamine depletion on sensitivity to rewarding and aversive stimuli in Parkinson's disease. *Neuropsychologia*, 51(5), 818–824. <https://doi.org/10.1016/j.neuropsychologia.2013.02.003>
- Martin, L. E., Cox, L. S., Brooks, W. M., & Savage, C. R. (2014). Winning and losing: differences in reward and punishment sensitivity between smokers and nonsmokers. *Brain and Behavior*, 4(6), 915–924. <https://doi.org/10.1002/brb3.285>

- Mata, R., Josef, A. K., Samanez-Larkin, G. R., & Hertwig, R. (2011). Age differences in risky choice: a meta-analysis. *Annals of the New York Academy of Sciences*, 1235, 18–29.
<https://doi.org/10.1111/j.1749-6632.2011.06200.x>
- Mattick, R. P., Breen, C., Kimber, J., & Davoli, M. (2009). Methadone maintenance therapy versus no opioid replacement therapy for opioid dependence. *Cochrane Database of Systematic Reviews*, 3. <https://doi.org/10.1002/14651858.CD002209.pub2>
- Mattson, C. L. (2021). Trends and Geographic Patterns in Drug and Synthetic Opioid Overdose Deaths—United States, 2013–2019. *MMWR. Morbidity and Mortality Weekly Report*, 70.
<https://doi.org/10.15585/mmwr.mm7006a4>
- Mell, T., Heekeren, H. R., Marschner, A., Wartenburger, I., Villringer, A., & Reischies, F. M. (2005). Effect of aging on stimulus-reward association learning. *Neuropsychologia*, 43(4), 554–563.
<https://doi.org/10.1016/j.neuropsychologia.2004.07.010>
- Mellis, A. M., & Bickel, W. K. (2020). Mechanical Turk data collection in addiction research: Utility, concerns and best practices. *Addiction (Abingdon, England)*, 115(10), 1960–1968.
<https://doi.org/10.1111/add.15032>
- Merikangas, K. R., Stolar, M., Stevens, D. E., Goulet, J., Preisig, M. A., Fenton, B., Zhang, H., O'Malley, S. S., & Rounsaville, B. J. (1998). Familial transmission of substance use disorders. *Archives of General Psychiatry*, 55(11), 973–979.
<https://doi.org/10.1001/archpsyc.55.11.973>
- Metin, B., Roeyers, H., Wiersema, J. R., van der Meere, J. J., Thompson, M., & Sonuga-Barke, E. (2013). ADHD performance reflects inefficient but not impulsive information processing: A diffusion model analysis. *Neuropsychology*, 27(2), 193–200.
<https://doi.org/10.1037/a0031533>
- Miletić, S., Boag, R. J., & Forstmann, B. U. (2020). Mutual benefits: Combining reinforcement learning with sequential sampling models. *Neuropsychologia*, 136, 107261.
<https://doi.org/10.1016/j.neuropsychologia.2019.107261>

- Mitchell, S. H. (1999). Measures of impulsivity in cigarette smokers and non-smokers. *Psychopharmacology*, *146*(4), 455–464.
- Moeller, S. J., & Paulus, M. P. (2018). Toward biomarkers of the addicted human brain: Using neuroimaging to predict relapse and sustained abstinence in substance use disorder. *Progress in Neuro-Psychopharmacology & Biological Psychiatry*, *80*(Pt B), 143–154.
<https://doi.org/10.1016/j.pnpbp.2017.03.003>
- Montague, P. R., Dolan, R. J., Friston, K. J., & Dayan, P. (2012). Computational psychiatry. *Trends in Cognitive Sciences*, *16*(1), 72–80. <https://doi.org/10.1016/j.tics.2011.11.018>
- Moreno Padilla, M., O'Halloran, L., Bennett, M., Cao, Z., & Whelan, R. (2017). Impulsivity and Reward Processing Endophenotypes in Youth Alcohol Misuse. *Current Addiction Reports*, *4*(4), 350–363. <https://doi.org/10.1007/s40429-017-0167-6>
- Morris, S. E., & Cuthbert, B. N. (2012). Research Domain Criteria: Cognitive systems, neural circuits, and dimensions of behavior. *Dialogues in Clinical Neuroscience*, *14*(1), 29–37.
- Moustafa, A. A., Kéri, S., Somlai, Z., Baldson, T., Frydecka, D., Misiak, B., & White, C. (2015). Drift diffusion model of reward and punishment learning in schizophrenia: Modeling and experimental data. *Behavioural Brain Research*, *291*, 147–154.
<https://doi.org/10.1016/j.bbr.2015.05.024>
- Muñoz, M. Á., Anllo-Vento, L., del Carmen Fernández, M., Montoya, P., & Vila, J. (2012). Modulation of the outcome-related negativity associated with nicotine abstinence. *Experimental and Clinical Psychopharmacology*, *20*(2), 151–160.
<https://doi.org/10.1037/a0025991>
- Myers, C. E., Rego, J., Haber, P., Morley, K., Beck, K. D., Hogarth, L., & Moustafa, A. A. (2017). Learning and generalization from reward and punishment in opioid addiction. *Behavioural Brain Research*, *317*, 122–131. <https://doi.org/10.1016/j.bbr.2016.09.033>
- Myers, C. E., Sheynin, J., Baldson, T., Luzzardo, A., Beck, K. D., Hogarth, L., Haber, P., & Moustafa, A. A. (2016). Probabilistic Reward- and Punishment-based Learning in Opioid Addiction:

Experimental and Computational Data. *Behavioural Brain Research*, 296, 240–248.

<https://doi.org/10.1016/j.bbr.2015.09.018>

Na, E., Jang, K.-M., & Kim, M.-S. (2019). An Event-Related Potential Study of Decision-Making and Feedback Utilization in Female College Students Who Binge Drink. *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.02606>

Nassar, M. R., Bruckner, R., Gold, J. I., Li, S.-C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, 7, 11609. <https://doi.org/10.1038/ncomms11609>

Nebe, S., Kroemer, N. B., Schad, D. J., Bernhardt, N., Sebold, M., Müller, D. K., Scholl, L., Kuitunen-Paul, S., Heinz, A., Rapp, M. A., Huys, Q. J. M., & Smolka, M. N. (2018). No association of goal-directed and habitual control with alcohol consumption in young adults: Alcohol use and learning. *Addiction Biology*, 23(1), 379–393. <https://doi.org/10.1111/adb.12490>

Nestor, L. J., McCabe, E., Jones, J., Clancy, L., & Garavan, H. (2018a). Shared and divergent neural reactivity to non-drug operant response outcomes in current smokers and ex-smokers. *Brain Research*, 1680, 54–61. <https://doi.org/10.1016/j.brainres.2017.12.003>

Nestor, L. J., McCabe, E., Jones, J., Clancy, L., & Garavan, H. (2018b). Smokers and ex-smokers have shared differences in the neural substrates for potential monetary gains and losses: Reward processing in smokers. *Addiction Biology*, 23(1), 369–378.

<https://doi.org/10.1111/adb.12484>

Neuhaus, A., Bajbouj, M., Kienast, T., Kalus, P., von Haebler, D., Winterer, G., & Gallinat, J. (2006). Persistent dysfunctional frontal lobe activation in former smokers. *Psychopharmacology*, 186(2), 191–200. <https://doi.org/10.1007/s00213-006-0366-7>

Nigbur, R., Ivanova, G., & Stürmer, B. (2011). Theta power as a marker for cognitive interference. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, 122(11), 2185–2194. <https://doi.org/10.1016/j.clinph.2011.03.030>

- Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3), 139–154. <https://doi.org/10.1016/j.jmp.2008.12.005>
- Notley, C., Gentry, S., Livingstone-Banks, J., Bauld, L., Perera, R., & Hartmann-Boyce, J. (2019). Incentives for smoking cessation. *Cochrane Database of Systematic Reviews*, 7. <https://doi.org/10.1002/14651858.CD004307.pub6>
- Nutt, D. J., King, L. A., & Phillips, L. D. (2010). Drug harms in the UK: A multicriteria decision analysis. *The Lancet*, 376(9752), 1558–1565. [https://doi.org/10.1016/S0140-6736\(10\)61462-6](https://doi.org/10.1016/S0140-6736(10)61462-6)
- O'Halloran, L., Pennie, B., Jollans, L., Kiiski, H., Vahey, N., Rai, L., Bradley, L., Lalor, R., & Whelan, R. (2018). A Combination of Impulsivity Subdomains Predict Alcohol Intoxication Frequency. *Alcoholism: Clinical and Experimental Research*, 42(8), 1530–1540. <https://doi.org/10.1111/acer.13779>
- Obermeyer, Z., & Emanuel, E. J. (2016). Predicting the Future—Big Data, Machine Learning, and Clinical Medicine. *The New England Journal of Medicine*, 375(13), 1216–1219. <https://doi.org/10.1056/NEJMp1606181>
- Palestro, J. J., Bahg, G., Sederberg, P. B., Lu, Z.-L., Steyvers, M., & Turner, B. M. (2018). A tutorial on joint models of neural and behavioral measures of cognition. *Journal of Mathematical Psychology*, 84, 20–48. <https://doi.org/10.1016/j.jmp.2018.03.003>
- Pandey, A. K., Kamarajan, C., Manz, N., Chorlian, D. B., Stimus, A., & Porjesz, B. (2016). Delta, theta, and alpha event-related oscillations in alcoholics during Go/NoGo task: Neurocognitive deficits in execution, inhibition, and attention processing. *Progress in Neuro-Psychopharmacology & Biological Psychiatry*, 65, 158–171. <https://doi.org/10.1016/j.pnpbp.2015.10.002>
- Parvaz, M. A., Konova, A. B., Proudfit, G. H., Dunning, J. P., Malaker, P., Moeller, S. J., Maloney, T., Alia-Klein, N., & Goldstein, R. Z. (2015). Impaired Neural Response to Negative Prediction Errors in Cocaine Addiction. *The Journal of Neuroscience*, 35(5), 1872–1879. <https://doi.org/10.1523/JNEUROSCI.2777-14.2015>

- Passetti, F., Clark, L., Davis, P., Mehta, M. A., White, S., Checinski, K., King, M., & Abou-Saleh, M. (2011). Risky decision-making predicts short-term outcome of community but not residential treatment for opiate addiction. Implications for case management. *Drug and Alcohol Dependence*, *118*(1), 12–18. <https://doi.org/10.1016/j.drugalcdep.2011.02.015>
- Pataky, T. C. (2012). One-dimensional statistical parametric mapping in Python. *Computer Methods in Biomechanics and Biomedical Engineering*, *15*(3), 295–301. <https://doi.org/10.1080/10255842.2010.527837>
- Patton, J.H. & Stanford, M.S. (1995). Factor structure of the Barratt impulsiveness scale. *Journal of clinical psychology*, *51*(6), 768-774.
- Paulus, M. P., Tapert, S. F., & Schuckit, M. A. (2005). Neural activation patterns of methamphetamine-dependent subjects during decision making predict relapse. *Archives of General Psychiatry*, *62*(7), 761–768. <https://doi.org/10.1001/archpsyc.62.7.761>
- Pedersen, M. L., Frank, M. J., & Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. *Psychonomic Bulletin & Review*, *24*(4), 1234–1251. <https://doi.org/10.3758/s13423-016-1199-y>
- Peechatka, A. L., & Janes, A. C. (2017). Association Between Reward Reactivity and Drug Use Severity is Substance Dependent: Preliminary Evidence From the Human Connectome Project. *Nicotine & Tobacco Research: Official Journal of the Society for Research on Nicotine and Tobacco*, *19*(6), 710–715. <https://doi.org/10.1093/ntr/ntw252>
- Peechatka, A. L., Whitton, A. E., Farmer, S. L., Pizzagalli, D. A., & Janes, A. C. (2015). Cigarette craving is associated with blunted reward processing in nicotine-dependent smokers. *Drug and Alcohol Dependence*, *155*, 202–207. <https://doi.org/10.1016/j.drugalcdep.2015.07.015>
- Perera, M. P. N., Bailey, N. W., Herring, S. E., & Fitzgerald, P. B. (2019). Electrophysiology of obsessive compulsive disorder: A systematic review of the electroencephalographic literature. *Journal of Anxiety Disorders*, *62*, 1–14. <https://doi.org/10.1016/j.janxdis.2018.11.001>

- Peters, L., & Soyka, M. (2019). Interrelationship of Opioid Dependence, Impaired Impulse Control, and Depressive Symptoms: An Open-Label Cross-Sectional Study of Patients in Maintenance Therapy. *Neuropsychobiology*, *77*(2), 73–82.
<https://doi.org/10.1159/000494697>
- Pfabigan, D. M., Alexopoulos, J., Bauer, H., & Sailer, U. (2011). Manipulation of feedback expectancy and valence induces negative and positive reward prediction error signals manifest in event-related brain potentials. *Psychophysiology*, *48*(5), 656–664.
<https://doi.org/10.1111/j.1469-8986.2010.01136.x>
- Pirrone, A., Dickinson, A., Gomez, R., Stafford, T., & Milne, E. (2017). Understanding perceptual judgment in autism spectrum disorder using the drift diffusion model. *Neuropsychology*, *31*(2), 173–180. <https://doi.org/10.1037/neu0000320>
- Poldrack, R. A., Feingold, F., Frank, M. J., Gleeson, P., de Hollander, G., Huys, Q. J. M., Love, B. C., Markiewicz, C. J., Moran, R., Ritter, P., Rogers, T. T., Turner, B. M., Yarkoni, T., Zhan, M., & Cohen, J. D. (2019). The Importance of Standards for Sharing of Computational Models and Data. *Computational Brain & Behavior*, *2*(3), 229–232. <https://doi.org/10.1007/s42113-019-00062-x>
- Potts, G. F., Bloom, E., Evans, D. E., & Drobles, D. J. (2014). Neural Reward and Punishment Sensitivity in Cigarette Smokers. *Drug and Alcohol Dependence*, *144*, 245–253.
<https://doi.org/10.1016/j.drugalcdep.2014.09.773>
- Powell, G., Jones, C. R. G., Hedge, C., Charman, T., Happé, F., Simonoff, E., & Sumner, P. (2019). Face processing in autism spectrum disorder re-evaluated through diffusion models. *Neuropsychology*, *33*(4), 445–461. <https://doi.org/10.1037/neu0000524>
- Prosser, J., Cohen, L. J., Steinfeld, M., Eisenberg, D., London, E. D., & Galyner, I. I. (2006). Neuropsychological functioning in opiate-dependent subjects receiving and following methadone maintenance treatment. *Drug and Alcohol Dependence*, *84*(3), 240–247.
<https://doi.org/10.1016/j.drugalcdep.2006.02.006>

- Radulescu, A., Daniel, R., & Niv, Y. (2016). The effects of aging on the interaction between reinforcement learning and attention. *Psychology and Aging, 31*(7), 747–757.
<https://doi.org/10.1037/pag0000112>
- Rai, L.A, O'Halloran, L., Jollans, L., Vahey, N., O'Brolchain, C., & Whelan, R. (2018). Individual differences in learning from probabilistic reward and punishment predicts smoking status. *Addictive Behaviors, 88*, 73-76.
- Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation, 20*(4), 873–922.
<https://doi.org/10.1162/neco.2008.12-06-420>
- Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science, 9*(5), 347–356. <https://doi.org/10.1111/1467-9280.00067>
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review, 9*(3), 438–481. <https://doi.org/10.3758/BF03196302>
- Ratcliff, R., Thapar, A., & McKoon, G. (2003). A diffusion model analysis of the effects of aging on brightness discrimination. *Perception & Psychophysics, 65*(4), 523–535.
<https://doi.org/10.3758/BF03194580>
- Ratcliff, R., Thapar, A., & McKoon, G. (2006). Aging, Practice, and Perceptual Tasks: A Diffusion Model Analysis. *Psychology and Aging, 21*(2), 353–371. <https://doi.org/10.1037/0882-7974.21.2.353>
- Ratcliff, R., Thapar, A., & McKoon, G. (2010). Individual Differences, Aging, and IQ in Two-Choice Tasks. *Cognitive Psychology, 60*(3), 127–157.
<https://doi.org/10.1016/j.cogpsych.2009.09.001>
- Ratcliff, R., Thapar, A., Gomez, P., & McKoon, G. (2004). A Diffusion Model Analysis of the Effects of Aging in the Lexical-Decision Task. *Psychology and Aging, 19*(2), 278–289.
<https://doi.org/10.1037/0882-7974.19.2.278>

- Redish, A. D. (2004). Addiction as a computational process gone awry. *Science (New York, N.Y.)*, 306(5703), 1944–1947. <https://doi.org/10.1126/science.1102384>
- Redish, A. D., Jensen, S., & Johnson, A. (2008). A unified framework for addiction: Vulnerabilities in the decision process. *The Behavioral and Brain Sciences*, 31(4), 415–487. <https://doi.org/10.1017/S0140525X0800472X>
- Reitsma, M. B., Fullman, N., Ng, M., Salama, J. S., Abajobir, A., Abate, K. H., Abbafati, C., Abera, S. F., Abraham, B., Abyu, G. Y., Adebisi, A. O., Al-Aly, Z., Aleman, A. V., Ali, R., Alkerwi, A. A., Allebeck, P., Al-Raddadi, R. M., Amare, A. T., Amberbir, A., ... Gakidou, E. (2017). Smoking prevalence and attributable disease burden in 195 countries and territories, 1990–2015: A systematic analysis from the Global Burden of Disease Study 2015. *The Lancet*, 389(10082), 1885–1906. [https://doi.org/10.1016/S0140-6736\(17\)30819-X](https://doi.org/10.1016/S0140-6736(17)30819-X)
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, 2, 64-99.
- Reynolds, B., Patak, M., Shroff, P., Penfold, R. B., Melanko, S., & Duhig, A. M. (2007). Laboratory and self-report assessments of impulsive behavior in adolescent daily smokers and nonsmokers. *Experimental and Clinical Psychopharmacology*, 15(3), 264–271. <https://doi.org/10.1037/1064-1297.15.3.264>
- Robbins, T. W., & Everitt, B. J. (1999). Drug addiction: Bad habits add up. *Nature*, 398(6728), 567–570. <https://doi.org/10.1038/19208>
- Robinson TE, Berridge KC. The neural basis of drug craving: an incentive-sensitization theory of addiction. *Brain Research. Brain Research Reviews*. 1993; 18(3):247–291.
- Robinson, O. J., & Chase, H. W. (2017). Learning and Choice in Mood Disorders: Searching for the Computational Parameters of Anhedonia. *Computational Psychiatry (Cambridge, Mass.)*, 1(1), 208–233. https://doi.org/10.1162/CPSY_a_00009

- Romanczuk-Seiferth, N., Koehler, S., Dreesen, C., Wüstenberg, T., & Heinz, A. (2015). Pathological gambling and alcohol dependence: Neural disturbances in reward and loss avoidance processing. *Addiction Biology, 20*(3), 557–569. <https://doi.org/10.1111/adb.12144>
- Rose, E. J., Ross, T. J., Salmeron, B. J., Lee, M., Shakleya, D. M., Huestis, M., Stein, E. A. (2012) Chronic exposure to nicotine is associated with reduced reward-related activity in the striatum but not the midbrain. *Biol Psychiatry 71*:206–213. DOI: 10.1016/j.biopsych.2011.09.013.
- Rosenthal, R. J., & Faris, S. B. (2019). The etymology and early history of ‘addiction.’ *Addiction Research & Theory, 27*(5), 437–449. <https://doi.org/10.1080/16066359.2018.1543412>
- Rossiter, S., Thompson, J., & Hester, R. (2012). Improving control over the impulse for reward: Sensitivity of harmful alcohol drinkers to delayed reward but not immediate punishment. *Drug and Alcohol Dependence, 125*(1–2), 89–94. <https://doi.org/10.1016/j.drugalcdep.2012.03.017>
- Rustemeier, M., Römling, J., Czybulka, C., Reymann, G., Daum, I., & Bellebaum, C. (2012). Learning from Positive and Negative Monetary Feedback in Patients with Alcohol Dependence. *Alcoholism: Clinical and Experimental Research, 36*(6), 560–573. <https://doi.org/10.1111/j.1530-0277.2011.01696.x>
- Saleme, D. M., Kluwe-Schiavon, B., Soliman, A., Misiak, B., Frydecka, D., & Moustafa, A. A. (2018). Factors underlying risk taking in heroin-dependent individuals: Feedback processing and environmental contingencies. *Behavioural Brain Research, 350*, 23–30. <https://doi.org/10.1016/j.bbr.2018.04.052>
- Salthouse, T. A. (1996). The processing-speed theory of adult age differences in cognition. *Psychological Review, 103*(3), 403–428.
- Sauseng, P., Griesmayr, B., Freunberger, R., & Klimesch, W. (2010). Control mechanisms in working memory: A possible function of EEG theta oscillations. *Neuroscience and Biobehavioral Reviews, 34*(7), 1015–1022. <https://doi.org/10.1016/j.neubiorev.2009.12.006>

- Schlienz, N. J., Hawk, L. W., & Rosch, K. S. (2013). The Effects of Acute Abstinence from Smoking and Performance-Based Rewards on Performance Monitoring. *Psychopharmacology*, 229(4), 701–711. <https://doi.org/10.1007/s00213-013-3131-8>
- Schubert, A.-L., Frischkorn, G. T., Hagemann, D., & Voss, A. (2016). Trait Characteristics of Diffusion Model Parameters. *Journal of Intelligence*, 4(3), 7. <https://doi.org/10.3390/jintelligence4030007>
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80(1), 1–27. <https://doi.org/10.1152/jn.1998.80.1.1>
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, 36(2), 241-263.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, 275(5306), 1593–1599.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593-1599.
- Schutte, I., Slagter, H. A., Collins, A. G. E., Frank, M. J., & Kenemans, J. L. (2017). Stimulus discriminability may bias value-based probabilistic learning. *PLOS ONE*, 12(5), e0176205. <https://doi.org/10.1371/journal.pone.0176205>
- Sebold, M., Nebe, S., Garbusow, M., Guggenmos, M., Schad, D. J., Beck, A., Kuitunen-Paul, S., Sommer, C., Frank, R., Neu, P., Zimmermann, U. S., Rapp, M. A., Smolka, M. N., Huys, Q. J. M., Schlagenhauf, F., & Heinz, A. (2017). When Habits Are Dangerous: Alcohol Expectancies and Habitual Decision Making Predict Relapse in Alcohol Dependence. *Biological Psychiatry*, 82(11), 847–856. <https://doi.org/10.1016/j.biopsych.2017.04.019>
- Seow, T. X. F., Benoit, E., Dempsey, C., Jennings, M., Maxwell, A., McDonough, M., & Gillan, C. M. (2019). Null results from a dimensional study of error-related negativity (ERN) and self-reported psychiatric symptoms. *BioRxiv*, 732594. <https://doi.org/10.1101/732594>

- Seow, Tricia X. F., O'Connell, R., & Gillan, C. M. (2020). Model-based learning deficits in compulsivity are linked to faulty representations of task structure. *BioRxiv*, 2020.06.11.147447. <https://doi.org/10.1101/2020.06.11.147447>
- Seth, P., Scholl, L., Rudd, R. A., & Bacon, S. (2018). Overdose Deaths Involving Opioids, Cocaine, and Psychostimulants—United States, 2015–2016. *MMWR. Morbidity and Mortality Weekly Report*, 67(12), 349–358. <https://doi.org/10.15585/mmwr.mm6712a1>
- Sher, K. J., Bartholow, B. D., & Wood, M. D. (2000). Personality and substance use disorders: A prospective study. *Journal of Consulting and Clinical Psychology*, 68(5), 818–829. <https://doi.org/10.1037/0022-006X.68.5.818>
- Shiner, T., Seymour, B., Wunderlich, K., Hill, C., Bhatia, K. P., Dayan, P., & Dolan, R. J. (2012). Dopamine and performance in a reinforcement learning task: evidence from Parkinson's disease. *Brain*, 135(6), 1871–1883. <https://doi.org/10.1093/brain/aws083>
- Sinha, R. (2011). New Findings on Biological Factors Predicting Addiction Relapse Vulnerability. *Current Psychiatry Reports*, 13(5), 398–405. <https://doi.org/10.1007/s11920-011-0224-0>
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York, London: D. Appleton-Century Company, Incorporated.
- Skinner, M. D., Aubin, H.-J., & Berlin, I. (2004). Impulsivity in smoking, nonsmoking, and ex-smoking alcoholics. *Addictive Behaviors*, 29(5), 973–978. <https://doi.org/10.1016/j.addbeh.2004.02.045>
- Smith, G. W., Farrell, M., Bunting, B. P., Houston, J. E., & Shevlin, M. (2011). Patterns of polydrug use in Great Britain: Findings from a national household population survey. *Drug and Alcohol Dependence*, 113(2–3), 222–228. <https://doi.org/10.1016/j.drugalcdep.2010.08.010>
- Smith, J. L., Mattick, R. P., Jamadar, S. D., & Iredale, J. M. (2014). Deficits in behavioural inhibition in substance abuse and addiction: A meta-analysis. *Drug and Alcohol Dependence*, 145, 1–33. Scopus. <https://doi.org/10.1016/j.drugalcdep.2014.08.009>

- Smyth, B. P., Barry, J., Keenan, E., & Ducray, K. (2010). Lapse and relapse following inpatient treatment of opiate dependence. *Irish Medical Journal*, *103*(6), 176–179.
- Sojitra, R. B., Lerner, I., Petok, J. R., & Gluck, M. A. (2018). Age affects reinforcement learning through dopamine-based learning imbalance and high decision noise—not through Parkinsonian mechanisms. *Neurobiology of Aging*, *68*, 102–113.
<https://doi.org/10.1016/j.neurobiolaging.2018.04.006>
- Soneji, S., Barrington-Trimis, J. L., Wills, T. A., Leventhal, A. M., Unger, J. B., Gibson, L. A., Yang, J., Primack, B. A., Andrews, J. A., Miech, R. A., Spindle, T. R., Dick, D. M., Eissenberg, T., Hornik, R. C., Dang, R., & Sargent, J. D. (2017). Association Between Initial Use of e-Cigarettes and Subsequent Cigarette Smoking Among Adolescents and Young Adults. *JAMA Pediatrics*, *171*(8), 788–797. <https://doi.org/10.1001/jamapediatrics.2017.1488>
- Spindle, T. R., Hiler, M. M., Cooke, M., Eissenberg, T., Kendler, K. S., & Dick, D. (2017). Electronic Cigarette Use and Uptake of Cigarette Smoking: A Longitudinal Examination of U.S. College Students. *Addictive Behaviors*, *67*, 66–72. <https://doi.org/10.1016/j.addbeh.2016.12.009>
- Sripada, C., & Weigard, A. (2021). Impaired Evidence Accumulation as a Transdiagnostic Vulnerability Factor in Psychopathology. *Frontiers in Psychiatry*, *12*.
<https://doi.org/10.3389/fpsy.2021.627179>
- Stead, L. F., Carroll, A. J., & Lancaster, T. (2017). Group behaviour therapy programmes for smoking cessation. *Cochrane Database of Systematic Reviews*, *3*.
<https://doi.org/10.1002/14651858.CD001007.pub3>
- Stein, J. S., Heckman, B. W., Pope, D. A., Perry, E. S., Fong, G. T., Cummings, K. M., & Bickel, W. K. (2018). Delay discounting and e-cigarette use: An investigation in current, former, and never cigarette smokers. *Drug and Alcohol Dependence*, *191*, 165–173. Scopus.
<https://doi.org/10.1016/j.drugalcdep.2018.06.034>
- Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2013). Validating the PVL-Delta model for the Iowa gambling task. *Frontiers in Psychology*, *4*. <https://doi.org/10.3389/fpsyg.2013.00898>

- Stewart, J. L., & May, A. C. (2016). Electrophysiology for addiction medicine: From methodology to conceptualization of reward deficits. *Progress in Brain Research*, 224, 67–84.
<https://doi.org/10.1016/bs.pbr.2015.07.013>
- Stewart, J. L., May, A. C., & Paulus, M. P. (2019). Bouncing back: Brain rehabilitation amid opioid and stimulant epidemics. *NeuroImage: Clinical*, 24, 102068.
<https://doi.org/10.1016/j.nicl.2019.102068>
- Stock, A.-K., Hoffmann, S., & Beste, C. (2017). Effects of binge drinking and hangover on response selection sub-processes—A study using EEG and drift diffusion modeling. *Addiction Biology*, 22(5), 1355–1365. <https://doi.org/10.1111/adb.12412>
- Strauss, G. P., Thaler, N. S., Matveeva, T. M., Vogel, S. J., Sutton, G. P., Lee, B. G., & Allen, D. N. (2015). Predicting psychosis across diagnostic boundaries: Behavioral and computational modeling evidence for impaired reinforcement learning in schizophrenia and bipolar disorder with a history of psychosis. *Journal of Abnormal Psychology*, 124(3), 697–708.
<https://doi.org/10.1037/abn0000039>
- Strickland, J. C., & Stoops, W. W. (2019). The use of crowdsourcing in addiction science research: Amazon Mechanical Turk. *Experimental and Clinical Psychopharmacology*, 27(1), 1–18.
<https://doi.org/10.1037/pha0000235>
- Su, H., Li, Z., Du, J., Jiang, H., Chen, Z., Sun, H., & Zhao, M. (2015). Predictors of heroin relapse: Personality traits, impulsivity, COMT gene Val158met polymorphism in a 5-year prospective study in Shanghai, China. *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics*, 168(8), 712–719. <https://doi.org/10.1002/ajmg.b.32376>
- Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning: an introduction. *IEEE Trans. Neural Netw.*, 9(5), 1054.
- Sutton, S. K., Van Rensburg, K. J., Jentink, K. G., Drobles, D. J., & Evans, D. E. (2016). Nicotine-induced cortical activation among nonsmokers with moderation by trait cognitive control. *Psychopharmacology*, 233(12), 2301–2308. <https://doi.org/10.1007/s00213-016-4276-z>

- Sweitzer, M. M., Geier, C. F., Addicott, M. A., Denlinger, R., Raiff, B. R., Dallery, J., McClernon, F. J., & Donny, E. C. (2016). Smoking Abstinence-Induced Changes in Resting State Functional Connectivity with Ventral Striatum Predict Lapse during a Quit Attempt. *Neuropsychopharmacology*, *41*(10), 2521–2529. Scopus. <https://doi.org/10.1038/npp.2016.56>
- Tanabe, J., Reynolds, J., Krmpotich, T., Claus, E., Thompson, L. L., Du, Y. P., & Banich, M. T. (2013). Reduced Neural Tracking of Prediction Error in Substance-Dependent Individuals. *American Journal of Psychiatry*, *170*(11), 1356–1363. <https://doi.org/10.1176/appi.ajp.2013.12091257>
- Thapar, A., Ratcliff, R., & McKoon, G. (2003). A Diffusion Model Analysis of the Effects of Aging on Letter Discrimination. *Psychology and Aging*, *18*(3), 415–429. <https://doi.org/10.1037/0882-7974.18.3.415>
- Torrens, M., Gilchrist, G., & Domingo-Salvany, A. (2011). Psychiatric comorbidity in illicit drug users: Substance-induced versus independent disorders. *Drug and Alcohol Dependence*, *113*(2), 147–156. <https://doi.org/10.1016/j.drugalcdep.2010.07.013>
- Tromp, D., Dufour, A., Lithfous, S., Pebayle, T., & Després, O. (2015). Episodic memory in normal aging and Alzheimer disease: Insights from imaging and behavioral studies. *Ageing Research Reviews*, *24*(Pt B), 232–262. <https://doi.org/10.1016/j.arr.2015.08.006>
- Turner, B.M., Palestro, J. J., Miletić, S., & Forstmann, B. U. (2019). Advances in techniques for imposing reciprocity in brain-behavior relations. *Neuroscience and Biobehavioral Reviews*, *102*, 327–336. Scopus. <https://doi.org/10.1016/j.neubiorev.2019.04.018>
- Turner, Brandon M., Rodriguez, C. A., Norcia, T. M., McClure, S. M., & Steyvers, M. (2016). Why more is better: Simultaneous modeling of EEG, fMRI, and behavioral data. *NeuroImage*, *128*, 96–115. <https://doi.org/10.1016/j.neuroimage.2015.12.030>
- Uhl, G. R., Koob, G. F., & Cable, J. (2019). The neurobiology of addiction. *Annals of the New York Academy of Sciences*, *1451*(1), 5–28. <https://doi.org/10.1111/nyas.13989>

- United Nations Office on Drugs and Crime, (2019). World Drug Report 2019. Available from:
<https://wdr.unodc.org/wdr2019/>
- Urošević, S., Halverson, T., Youngstrom, E. A., & Luciana, M. (2018). Probabilistic reinforcement learning abnormalities and their correlates in adolescent bipolar disorders. *Journal of Abnormal Psychology, 127*(8), 807–817. <https://doi.org/10.1037/abn0000388>
- van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E.-J. (2012). A diffusion model decomposition of the effects of alcohol on perceptual decision making. *Psychopharmacology, 219*(4), 1017–1025. <https://doi.org/10.1007/s00213-011-2435-9>
- Vanderschuren, L. J., Minnaard, A. M., Smeets, J. A., & Lesscher, H. M. (2017). Punishment models of addictive behavior. *Current Opinion in Behavioral Sciences, 13*, 77–84.
<https://doi.org/10.1016/j.cobeha.2016.10.007>
- Vassileva, J., Paxton, J., Moeller, F. G., Wilson, M. J., Bozgunov, K., Martin, E. M., Gonzalez, R., & Vasilev, G. (2014). Heroin and amphetamine users display opposite relationships between trait and neurobehavioral dimensions of impulsivity. *Addictive Behaviors, 39*(3), 652–659.
<https://doi.org/10.1016/j.addbeh.2013.11.020>
- Villafuerte, S., Heitzeg, M. M., Foley, S., Yau, W.-Y. W., Majczenko, K., Zubieta, J.-K., Zucker, R. A., & Burmeister, M. (2012). Impulsiveness and insula activation during reward anticipation are associated with genetic variants in GABRA2 in a family sample enriched for alcoholism. *Molecular Psychiatry, 17*(5), 511–519. <https://doi.org/10.1038/mp.2011.33>
- Volkow, N. D., Koob, G. F., & McLellan, A. T. (2016). Neurobiologic Advances from the Brain Disease Model of Addiction. *New England Journal of Medicine, 374*(4), 363–371.
<https://doi.org/10.1056/NEJMra1511480>
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior* (pp. xviii, 625). Princeton University Press.
- Vos, T., Abajobir, A. A., Abate, K. H., Abbafati, C., Abbas, K. M., Abd-Allah, F., Abdulkader, R. S., Abdulle, A. M., Abebo, T. A., Abera, S. F., Aboyans, V., Abu-Raddad, L. J., Ackerman, I. N.,

- Adamu, A. A., Adetokunboh, O., Afarideh, M., Afshin, A., Agarwal, S. K., Aggarwal, R., ... Murray, C. J. L. (2017). Global, regional, and national incidence, prevalence, and years lived with disability for 328 diseases and injuries for 195 countries, 1990–2016: A systematic analysis for the Global Burden of Disease Study 2016. *The Lancet*, *390*(10100), 1211–1259. [https://doi.org/10.1016/S0140-6736\(17\)32154-2](https://doi.org/10.1016/S0140-6736(17)32154-2)
- Wahlstrom, L. C. (n.d.). *Feedback-Related Negativity, Decision-Making, and College Binge Drinking*. 118.
- Waltz, J. A., Frank, M. J., Wiecki, T. V., & Gold, J. M. (2011). Altered probabilistic learning and response biases in schizophrenia: Behavioral evidence and neurocomputational modeling. *Neuropsychology*, *25*(1), 86–97. <https://doi.org/10.1037/a0020882>
- Warner, M., Chen, L. H., Makuc, D. M., Anderson, R. N., & Miniño, A. M. (2011). Drug poisoning deaths in the United States, 1980-2008. *NCHS Data Brief*, *81*, 1–8.
- Weigard, A. S., Brislin, S. J., Cope, L. M., Hardee, J. E., Martz, M. E., Ly, A., Zucker, R. A., Sripada, C., & Heitzeg, M. M. (2021). Evidence accumulation and associated error-related brain activity as computationally-informed prospective predictors of substance use in emerging adulthood. *BioRxiv*, 2020.03.06.981035. <https://doi.org/10.1101/2020.03.06.981035>
- Wennerhold, L., & Friese, M. (2020). Why Self-Report Measures of Self-Control and Inhibition Tasks Do Not Substantially Correlate. *Collabra: Psychology*, *6*(9). <https://doi.org/10.1525/collabra.276>
- Whelan, R., & Garavan, H. (2014). When optimism hurts: Inflated predictions in psychiatric neuroimaging. *Biological Psychiatry*, *75*(9), 746–748. <https://doi.org/10.1016/j.biopsych.2013.05.014>
- Whelan, R., Watts, R., Orr, C. A., Althoff, R. R., Artiges, E., Banaschewski, T., G.J. Barker, A.L. Bokde, C. Buchel, F.M. Carvalho, P.J. Conrod, H. Flor, M. Fauth-Buhler, V. Frouin, J. Gallinat, G. Gan, P. Gowland, A. Heinz, B. Ittermann, C. Lawrence, K. Mann, J.L. Martinot, F. Nees, N. Ortiz, M.L. Paillere-Martinot, T. Paus, Z. Pausova, M. Rietschel, T.W. Robbins, Smolka, M.N.,

- Strohle, A., Schumann, G., & Garavan, H. (2014). Neuropsychosocial profiles of current and future adolescent alcohol misusers. *Nature*, 512(7513), 185-189
- White, C. N., Ratcliff, R., Vasey, M. W., & McKoon, G. (2010). Using diffusion models to understand clinical disorders. *Journal of Mathematical Psychology*, 54(1), 39–52.
<https://doi.org/10.1016/j.jmp.2010.01.004>
- Whiteford, H. A., Degenhardt, L., Rehm, J., Baxter, A. J., Ferrari, A. J., Erskine, H. E., Charlson, F. J., Norman, R. E., Flaxman, A. D., Johns, N., Burstein, R., Murray, C. J., & Vos, T. (2013). Global burden of disease attributable to mental and substance use disorders: Findings from the Global Burden of Disease Study 2010. *The Lancet*, 382(9904), 1575–1586.
[https://doi.org/10.1016/S0140-6736\(13\)61611-6](https://doi.org/10.1016/S0140-6736(13)61611-6)
- Whitmer, A. J., Frank, M. J., & Gotlib, I. H. (2012). Sensitivity to reward and punishment in major depressive disorder: effects of rumination and of single versus multiple experiences. *Cognition & Emotion*, 26(8), 1475–1485. <https://doi.org/10.1080/02699931.2012.682973>
- Whittaker, R., McRobbie, H., Bullen, C., Rodgers, A., Gu, Y., & Dobson, R. (2019). Mobile phone text messaging and app-based interventions for smoking cessation. *Cochrane Database of Systematic Reviews*, 10. <https://doi.org/10.1002/14651858.CD006611.pub5>
- Wiecki, T. V., Antoniades, C. A., Stevenson, A., Kennard, C., Borowsky, B., Owen, G., Leavitt, B., Roos, R., Durr, A., Tabrizi, S. J., & Frank, M. J. (2016). A Computational Cognitive Biomarker for Early-Stage Huntington’s Disease. *PLOS ONE*, 11(2), e0148409.
<https://doi.org/10.1371/journal.pone.0148409>
- Wiecki, T. V., Antoniades, C. A., Stevenson, A., Kennard, C., Borowsky, B., Owen, G., Leavitt, B., Roos, R., Durr, A., Tabrizi, S. J., & Frank, M. J. (2016). A Computational Cognitive Biomarker for Early-Stage Huntington’s Disease. *PLOS ONE*, 11(2), e0148409.
<https://doi.org/10.1371/journal.pone.0148409>

- Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics, 7*.
<https://doi.org/10.3389/fninf.2013.00014>
- Wiehler, A., & Peters, J. (2020). Diffusion modeling reveals reinforcement learning impairments in gambling disorder that are linked to attenuated ventromedial prefrontal cortex value representations. *BioRxiv*, 2020.06.03.131359. <https://doi.org/10.1101/2020.06.03.131359>
- Williams, J. K. (2009). Reinforcement Learning of Optimal Controls. In S. E. Haupt, A. Pasini, & C. Marzban (Eds.), *Artificial Intelligence Methods in the Environmental Sciences* (pp. 297–327). Springer Netherlands. https://doi.org/10.1007/978-1-4020-9119-3_15
- Wilson, R. C., & Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *ELife, 8*, e49547. <https://doi.org/10.7554/eLife.49547>
- World Drug Report (2018). Available from: <https://www.unodc.org/wdr2018/>
- World Health Organisation (2010). Addiction to Nicotine. Available from:
https://www.who.int/tobacco/publications/gender/en_tfi_gender_women_addiction_nicotine.pdf
- World Health Organisation (2018). Global status report on alcohol and health: 2018. Available from:
<https://www.who.int/publications/i/item/9789241565639>
- World Health Organization (2020). International statistical classification of diseases and related health problems (11th ed.). <https://icd.who.int/>
- Wrase, J., Schlagenhaut, F., Kienast, T., Wüstenberg, T., Bermanpohl, F., Kahnt, T., Beck, A., Ströhle, A., Juckel, G., Knutson, B., & Heinz, A. (2007). Dysfunction of reward processing correlates with alcohol craving in detoxified alcoholics. *NeuroImage, 35*(2), 787–794.
<https://doi.org/10.1016/j.neuroimage.2006.11.043>
- Wrase, J., Schlagenhaut, F., Kienast, T., Wüstenberg, T., Bermanpohl, F., Kahnt, T., Beck, A., Ströhle, A., Juckel, G., Knutson, B., & Heinz, A. (2007). Dysfunction of reward processing correlates with

alcohol craving in detoxified alcoholics. *NeuroImage*, 35(2), 787–794.

<https://doi.org/10.1016/j.neuroimage.2006.11.043>

Wu, Y., & Zhou, X. (2009). The P300 and reward valence, magnitude, and expectancy in outcome evaluation. *Brain Research*, 1286, 114–122. <https://doi.org/10.1016/j.brainres.2009.06.032>

Wu, Y., & Zhou, X. (2009). The P300 and reward valence, magnitude, and expectancy in outcome evaluation. *Brain Research*, 1286, 114–122. <https://doi.org/10.1016/j.brainres.2009.06.032>

Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 12(6), 1100–1122. <https://doi.org/10.1177/1745691617693393>

Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 12(6), 1100–1122. <https://doi.org/10.1177/1745691617693393>

Zhao, Q., Li, H., Hu, B., Wu, H., & Liu, Q. (2017). Abstinent Heroin Addicts Tend to Take Risks: ERP and Source Localization. *Frontiers in Neuroscience*, 11.

<https://doi.org/10.3389/fnins.2017.00681>

Zhao, Q., Li, H., Hu, B., Wu, H., & Liu, Q. (2017). Abstinent Heroin Addicts Tend to Take Risks: ERP and Source Localization. *Frontiers in Neuroscience*, 11.

<https://doi.org/10.3389/fnins.2017.00681>

Zhu, Y., Evans, E., Mooney, L. J., Saxon, A., Kelleghan, A., Yoo, C., & Hser, Y.-I. (2018a). Correlates of long-term opioid abstinence after randomization to methadone versus buprenorphine/naloxone in a multi-site trial. *Journal of Neuroimmune Pharmacology : The Official Journal of the Society on NeuroImmune Pharmacology*, 13(4), 488–497.

<https://doi.org/10.1007/s11481-018-9801-x>

Zhu, Y., Evans, E., Mooney, L. J., Saxon, A., Kelleghan, A., Yoo, C., & Hser, Y.-I. (2018b). Correlates of long-term opioid abstinence after randomization to methadone versus buprenorphine/naloxone in a multi-site trial. *Journal of Neuroimmune Pharmacology : The*

Official Journal of the Society on NeuroImmune Pharmacology, 13(4), 488–497.

<https://doi.org/10.1007/s11481-018-9801-x>

Ziegler, S., Pedersen, M. L., Mowinckel, A. M., & Biele, G. (2016). Modelling ADHD: A review of ADHD theories through their predictions for computational models of decision-making and reinforcement learning. *Neuroscience & Biobehavioral Reviews*, 71, 633–656.

<https://doi.org/10.1016/j.neubiorev.2016.09.002>

Zou, H. & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301-320.