

From Labs to Real-World: Developing Smartphone-Based Methodologies for Enhanced Phenotyping of Human Decision-Making in Clinical Settings

Lili Zhang

B.Eng., M.Eng.

A dissertation submitted in fulfilment of the requirements for the award of

Doctor of Philosophy (Ph.D.)

to the



School of Computing

Dublin City University

Supervisor: Prof. Tomás Ward

July 2023

Declaration

I hereby certify that this material, which I now submit for assessment on the programme of study leading to the award of Doctor of Philosophy is entirely my own work, that I have exercised reasonable care to ensure that the work is original, and does not to the best of my knowledge breach any law of copyright, and has not been taken from the work of others save and to the extent that such work has been cited and acknowledged within the text of my work.

Signed: _____ (Candidate) ID No: 18214884 Date: July 21, 2023

Acknowledgements

First and foremost, I would like to express my deepest appreciation to my supervisor Prof. Tomás E. Ward. I would like to thank him for giving me the opportunity to take this PhD position four years ago, which was in fact a major turn point of my life. This opportunity allowed me to switch my research to Computational Phenotyping of human Decision-making, which was a completely new but very appealing area to me. It was also this opportunity which made my dream of studying abroad become reality. Additionally, this PhD opportunity allowed me to come to Ireland to reunion with my boyfriend, who had been doing his PhD in Ireland for one year back then. I would also like to thank Tomás for his constant guidance, encouragement, caring and support during my PhD. He always encourages me when I was frustrated or when my paper was rejected in the first place. He was always sincerely happy for me and generously offers his praises to me whenever I had some progress and achievements, which were so helpful for building my confidence. He barely pushes me on my work like a boss, instead he provides valuable suggestions and encourages me to learn new skills and try out new ideas. I am really lucky to have Tomás as my supervisor throughout my PhD. I am also deeply grateful to my master supervisor Prof. Libao Deng for his understanding and tolerance. I started my PhD while I was still doing my master degree, which would not be possible without his support.

I would also like to thank all of my wonderful friends and colleagues in Dublin City University as well as the entire Insight DCU family for their constant support. I am honored to be in the same group with you, Long Cheng, Zhengwei Wang, Jose Juan Dominguez Veiga, Eoin Brophy, Greta Monacelli, and Himanshu Vashisht. I would like to convey my sincere appreciation to Himanshu who was involved in most of my PhD projects. Thanks so much for your time and patience. I really enjoy talking and discussing with you during my PhD. I would like to thank our collaborators, Dr. Brian Slattery, Dr. Veronic Bruno and Emma Meade. I could not finish my projects smoothly without your help. I would like to thank Liting Zhou, Lola and all of our lovely friends in Dublin for your company. It is also nice to know my other two friends, Xiao Huang and Bo Zhang, after moving to Cork. I enjoyed every moment we shared together. Thanks for making me not feeling lonely at abroad. I spent nine months working as an intern in the group of Advanced Analytics at AIB. It was in fact the first time of my life working in the industry. I would like to express my gratitude to Paul Hunter, the head of the group, for giving me the opportunity to add this wonderful working experience to my CV. I would also like to thank all members in the group, particularly Cathal O'Donovan and Eamon McEntee, for guiding me and producing a warm working atmosphere. In the last four years, I was supported by Insight Centre for Data Analytics which is sponsored by Science Foundation Ireland under Grant Number *SFI/12/RC/2289_P2*. Their generous support is also highly acknowledged.

I would also wish to extend my sincere thanks to my aunt and uncle, who treat me like

their daughter. Thanks for your love, caring, and support. Last but most importantly, thank you Yaoyi. You were the initial reason why I came to Ireland and I never regretted for this decision. It is a great pleasant and an honor to be part of your life. Thanks for loving me, supporting me, and tolerating my bad tempers.

List of Publications

The following are journal/book chapter papers that have been submitted/published during the course of my PhD.

- **Zhang, L.**, Vashisht, H., Nethra, A., Slattery, B., & Ward, T. (2022). Differences in Learning and Persistency Characterizing Behavior in Chronic Pain for the Iowa Gambling Task: Web-Based Laboratory-in-the-Field Study. *Journal of Medical Internet Research*, 24(4), e26307.
- **Zhang, L.**, Monacelli, G., Vashisht, H., Schlee, W., Langguth, B., & Ward, T. (2022). The Effects of Tinnitus in Probabilistic Learning Tasks: Protocol for an Ecological Momentary Assessment Study. *JMIR Research Protocols*, 11(11), e36583.
- **Zhang, L.**, Vashisht, H., Totev, A., Trinh, N., & Ward, T. (2022). A Comparison of Distributed Machine Learning Methods for the Support of "Many Labs" Collaborations in Computational Modelling of Decision Making. *Frontiers in Psychology*, 13.
- McCarthy, M., **Zhang, L.**, Monacelli, G., & Ward, T. (2021). Using Methods From Computational Decision-making to Predict Nonadherence to Fitness Goals: Protocol for an Observational Study. *JMIR research protocols*, 10(11), e29758.
- Monacelli, G., **Zhang, L.**, Schlee, W., Langguth, B., Ward, T. E., & Murphy, T. B. (2022). Adaptive data collection for intra-individual studies affected by adherence. arXiv preprint arXiv:2207.12331. (**accepted by Biometrical Journal**)

The following are conference papers that have been published during the course of my PhD.

- **Zhang, L.**, Mukherjee, R., Wadhai, P., Muehlhausen, W., & Ward, T. (2022). Computational Phenotyping of Decision-making over Voice Interfaces. 30th Irish Conference on Artificial Intelligence and Cognitive Science (AICS).
- **Zhang, L.**, Boyle, D. O., Vashisht, H., & Ward, T. (2022). Belief in Conspiracy Theories is associated with decreased adaptive learning to contingency volatility. 22nd Conference of the European Society for Cognitive Psychology (ESCOPE).
- Betzhold, S., Donnelly J., **Zhang L.**, Ridge K., Ward T., & Conlon, N. (2022). STOP CSUA: Activity, Mood and Sleep in Chronic Spontaneous Urticaria and Angioedema. The 6th GA²LEN Global Urticaria Forum.
- Muehlhausen, W., **Zhang, L.**, Smith, L., & Ward, T. (2019). PNS230 Using Machine Learning for Signal Detection in Real-world Data From Wristworn Wearable Devices to Identify Fraudulent Behavior. *Value in Health*, 22, S325.

-
- **Zhang, L.,** Vashisht, H., Nethra, A., Slattery, B., & Ward, T. (2021). Altered Executive Functions Associated With Chronic Pain is Age-Independent. *Biological Psychiatry*, 89(9), S310.
 - **Zhang, L.,** & Ward, T. (2020). The Application of an Outcome-Representation Learning Model for Characterization of the Decision-Making Strategies Used by Younger and Older People. *Biological Psychiatry*, 87(9), S229-S230.

Contents

Declaration	i
Acknowledgments	iii
List of Publications	v
List of Figures	xiii
List of Tables	xv
Abstract	xvii
1 Introduction	1
1.1 Overview	1
1.2 Research Objectives	4
1.3 Research Questions	6
1.4 Organization of the Thesis	9
2 Methodology and Theoretical Foundations	11
2.1 Methodology 1: Lab-in-the-field Experiments	12
2.1.1 Decision-making Paradigms	12
2.1.2 Laboratory Experiments	12
2.1.3 Lab-in-the-field Experiments	13
2.2 Methodology 2: EMA	14
2.3 Methodology Implementation	15
2.3.1 Historical Methods	16
2.3.2 Smartphones	16
2.3.3 Data Quality	18
2.4 Economic Theories of Decision-making	19
2.4.1 Economic Theories of Choice	19
2.4.2 Expected Value Theory	20
2.4.3 Expected Utility Theory	21
2.4.4 Prospect Theory	23
2.4.5 Neural Representations of Stimulus Values	25
2.5 Reinforcement Learning Theories	27
2.5.1 Decisions from Experience	27
2.5.2 Rescorla-Wager Model	27
2.5.3 Transforming Values into Choice	30
2.5.4 Responses of Dopamine Neurons to Prediction Error Signal	32

2.6	Deep Learning Models to Predict Human Behavior	35
2.7	Conclusion	37
3	Tools for Data Collection and Analysis	39
3.1	The Web-based Decision-making Paradigms	39
3.1.1	Introduction	39
3.1.2	The Web-based Iowa Gambling Task	40
3.1.3	The Web-based ALT	42
3.2	EMA Platform: AthenaCX	45
3.2.1	Introduction	45
3.2.2	E-consent and Survey Design	45
3.2.3	Sampling Schemes	47
3.2.4	Push Notifications	49
3.2.5	Offline Data Collection	50
3.2.6	Link with External Web-based Tasks	50
3.2.7	Participants Onboarding	51
3.2.8	Data Security	51
3.3	Workflow of Theory-Driven Computational Modeling	52
3.3.1	Introduction	52
3.3.2	Model Simulation	53
3.3.3	Parameter Estimation	54
3.3.4	Model Comparison	56
3.3.5	Posterior Predictive Check	58
3.3.6	Parameter Recovery	59
3.4	Conclusion	60
4	Differences in Reward and Persistency Characterize Behavior in Chronic Pain for the IGT: A Lab-in-the-field Study	61
4.1	Introduction	62
4.2	Methods	64
4.2.1	Recruitment	64
4.2.2	Assessment of Pain Experience	65
4.2.3	The Web-based IGT	66
4.2.4	Standard Behavioral Data Analysis	66
4.2.5	Computational Modelling Analysis	67
4.3	Results	71
4.3.1	Self-report Analysis	71
4.3.2	Standard Behavioral Data Analysis	73
4.3.3	Computational Modelling Analysis for the IGT	75
4.4	Discussion	80
4.5	Conclusion	82
5	Impaired Adaptation of Learning to Contingency Volatility is Attributed to the Depression and Anxiety Symptoms of Parkinson’s Disease: A Lab-in-the-field Study	85
5.1	Introduction	86
5.2	Methods	89
5.2.1	Recruitment	89
5.2.2	Study Design	90

5.2.3	Model-agnostic Behavioral Analysis	92
5.2.4	Computational Modelling Analysis	93
5.3	Results	98
5.3.1	Demographic and Clinical Data	98
5.3.2	Model-agnostic Analysis	99
5.3.3	Computational Modelling Analysis	100
5.4	Discussion	104
5.5	Conclusion	108
6	The Effects of Tinnitus in Probabilistic Learning Tasks: an EMA study	109
6.1	Introduction	110
6.1.1	Background	110
6.1.2	Tinnitus and Cognitive Impairments	110
6.1.3	The Relationship between Tinnitus, Cognition and Psychological Disorders	111
6.1.4	Inclusion of Tinnitus Symptom Dynamics	112
6.2	Methods	113
6.2.1	Recruitment	113
6.2.2	Study Design	114
6.2.3	Study App	116
6.2.4	The ESIT-SQ Measure	117
6.2.5	The MTQ Measure	117
6.2.6	The STAI Measure	117
6.2.7	The Major Depression Inventory	118
6.2.8	The EMA Survey of Tinnitus Symptom and Emotional Status	118
6.2.9	The Web-based Aversive Learning Task	120
6.2.10	Compliance Analysis of the Longitudinal Part of the Experiment	120
6.2.11	Model-agnostic Behavioral Statistics	121
6.2.12	Group-level Computational Phenotypes across Participants and Sessions	121
6.2.13	Task performance Variability and Impact of EMA Variables	122
6.3	Results	123
6.3.1	Demographic Data	123
6.3.2	Compliance Analysis	123
6.3.3	Model-agnostic Behavioral Statistics	124
6.3.4	Group-level Computational Phenotypes across Participants and Sessions	126
6.3.5	Task Performance Variability and Impacts of EMA variables	127
6.4	Discussion	129
6.5	Conclusion	131
7	A Comparison of Distributed Machine Learning Methods for the Support of Large-scale Computational Phenotyping of Decision-making	133
7.1	Introduction	134
7.2	Methods	137
7.2.1	The IGT dataset	137
7.2.2	RNN Model for the IGT	138
7.2.3	The Architectures of the Four Training Strategies	141

7.3	Results	142
7.3.1	Experimental Settings	142
7.3.2	The Need for More Numerous and Diverse Data	144
7.3.3	The Benefits of Training Collaborative Models	147
7.3.4	On-policy Simulation	147
7.3.5	Off-policy Simulation	153
7.4	Discussion	155
7.5	Conclusion	158
8	Conclusion	159
8.1	Summary	159
8.2	Limitations	163
8.3	Future Work	165
8.3.1	Further Applications of the Lab-in-the-Field Methodology	165
8.3.2	Extension of the EMA Method	166
8.3.3	Scaling the Experiment with the Voice-based Devices	167
8.3.4	Summary	168

List of Abbreviations

RL	Reinforcement Learning
ML	Machine Learning
EMA	Ecological Momentary Assessment
BYOD	“Bring Your Own Device”
ePRO	electronic Patient Reported Outcomes
PDA	Personal Data Assistant
SMS	Short Message Service
SV	Stimulus Values
BOLD	blood-oxygenation-level-dependent
fMRI	functional Magnetic Resonance Imaging
EEG	electroencephalography
MEG	magnetoencephalography
vmPFC	ventromedial prefrontal cortex
IPFC	lateral prefrontal cortex
OFC	orbitofrontal cortex
ACC	anterior cingulate cortex
SEF	supplementary eye field
RW	Rescorla-Wagner
DL	Deep Learning
RNN	Recurrent Neural Network
IGT	Iowa Gambling Task
ALT	Aversive Learning Task
DS	Design Studio
VRS	Verbal Response Scale

VAS	Visual Analogue Scale
NRS	Numeric Response Scale
SSL	Secure Sockets Layer
MLE	Maximum Likelihood Estimation
HBM	Hierarchical Bayesian Modeling
MCMC	Markov Chain Monte Carlo
AIC	Akaike Information Criterion
DIC	Deviance Information Criterion
WAIC	Widely Applicable Information Criterion
LOOCV	leave-one-out cross-validation
PSIS	Pareto-smoothed Importance Sampling
PPC	posterior predictive check
BPI	Brief Pain Inventory
EVL	Expectancy-Valence Learning
PVL-Delta	Prospect Valence Learning model with Delta
PVL-decay	Prospect Valence Learning model with decay
VPP	Values-Plus-Perseverance
ORL	Outcome-Representation Learning
PDQ	Parkinson Disease Questionnaire
MDS-UPDRS	MDS-Unified Parkinson's Disease Rating Scale
HDRS	Hamilton Depression Rating Scale
HAM-A	Hamilton Scale for Anxiety
KPPS	King's Parkinson's Disease Pain Scale
SAS	Starkstein Apathy Scale
LED	Levodopa Equivalent Dose
ESIT-SQ	European School for Interdisciplinary Tinnitus Research-Screening Questionnaire
STAI	State-Trait Anxiety Inventory
MDI	Major Depression Inventory

MTQ	Mini Tinnitus Questionnaire
LME	Linear Mixed Effects
ICC	Intraclass Correlation Coefficient
FL	Federated Learning
CL	Centralized Learning
HIPAA	Health Insurance Portability and Accountability Act
GDPR	General Data Protection Regulation
IL	Incremental Learning
CIL	Cyclic Incremental Learning

List of Figures

2.1	A representative value function proposed by Tversky and Kahneman [62].	24
2.2	The influence of recent outcomes on choice for three different learning rates ([84]). Red represents a participant with a relatively high learning rate of 0.9. There is a strong influence of the outcome on the recent previous trial ($n - 1$), and a weaker influence on the outcome received two trials back ($n - 2$) but virtually no influence on earlier outcomes. A learner with smaller learning rates, here shown for 0.5 (grey) or 0.2 (blue), shows increasingly longer-lasting influences of outcomes received on trials further back from the current trial.	30
2.3	The impacts of inverse temperature on the steepness of the softmax function.	31
2.4	The three stages of the theoretical accounts of experience-based decision-making process proposed in [80]. (i) Calculate the values of the available stimuli (here a, b, c) in the current situation. (ii) Choose and execute one by comparing the expected rewards of each option. (iii) Finally, learn from the reward prediction error to improve future decisions. Numbers indicate the predicted action values, the obtained reward, and the resulting prediction error.	32
2.5	Rasters and histograms from a primate dopaminergic neuron in a classical conditional experiment in [74], illustrating the responses to the conditioned stimuli and reward at various reward probabilities. The five panels correspond to five cues trained with stochastic reinforcement at different probabilities, increasing from top to bottom. Dopamine neurons responses increase with the size of the prediction error.	34
2.6	Structure of the Recurrent Neural Network (RNN) model proposed in [92]. The LSTM layer circled in red dashed line receives the previous action and reward as inputs, and is connected to a softmax layer (inside the black rectangle) which outputs the probability of selecting each action on the next trial. The LSTM layer is composed of a set of LSTM cells (shown by blue circles), which are connected to each other (shown by green arrows). The output of the cells (denoted by h_t^i for cell i at time t) are connected to a softmax layer using a set of connections shown by black lines. The free parameters of the model are denoted by Θ , and $(\Theta; RNN)$ is a metric which represents how well the model fits subjects' data and is used to adjust the parameters of the model using the MLE as the network learns how humans learn.	36

2.7 The structure of the autoencoder-based framework proposed in [94]. The model comprises an encoder (red) and a decoder (blue). The encoder is the composition of an RNN and a series of fully-connected layers. The encoder maps each whole input sequence (in the rectangles on the left) into a point in the latent space (depicted by (z_1, z_2) -coordinates in the middle) based on the final state of the RNN. The latent representation for each input sequence is then fed into the decoder. The decoder generates the weight parameters of an RNN (called the learning network here) and is shown by the dotted lines on the right side. The learning network is the reconstruction of the closed-loop dynamical process of the subject which generated the corresponding input sequence based on experience. This learning network has similar structure shown in Figure 2.6. 37

3.1 The screenshot of the web-based IGT. 41

3.2 The screenshots of the web-based ALT. In the first trial (panel A), the subject chose the blue leprechaun, which did not steal gold coins from the subject, thus the feedback is "Good choice". In the second trial (panel B), the subject chose the red leprechaun and it stole 36 gold coins and ran away. 43

3.3 Outcome probabilities across the course of the ALT. The task comprised two blocks. In the stable block, one leprechaun (for example, the blue one) had a 75% probability of resulting in loss of gold coins; the other shape had a 25% probability of resulting in coin loss. In the volatile block, the probability that the choice of a given leprechaun would result in coin loss switched every 30 trials between 80% and 20%. 43

3.4 The home page screenshot of the DS. 46

3.5 Question types provided in AthenaCX. 46

3.6 Survey types provided in AthenaCX. Inside the panel on the right side of the survey card, the alarm clock is the time-based survey. Researchers can use this to create schedules based on specified timing patterns, e.g. a daily or weekly survey. The triangle is the Self-Initiated survey. This survey type will always be available to the participant during the time period that it has been configured to be available. The third sharing icon is a remote trigger survey. This survey type can be activated or scheduled directly from the dashboard at any point once the participant has enrolled. 48

4.1 A graphical representation of the conceptual framework of the measurement. 65

4.2 Pain severity plotted against pain interference (left); activity sub-dimension plotted against affective sub-dimension (right). Healthy controls ($N = 25$) are plotted in red and people living with chronic pain ($N = 20$) are plotted in blue. The r values were calculated between the paired pain measures for the whole sample ($*P < 0.5, **P < 0.01, ***P < 0.001$). 72

4.3 Pain severity plotted against total amount of gain by the end of the task. Healthy controls ($N = 25$) are plotted in red and people living with chronic pain ($N = 20$) are plotted in blue. The r value was calculated between the pain measure and the task performance measure for the whole sample ($*P < 0.5, **P < 0.01, ***P < 0.001$). 72

4.4	Mean proportion of choices from each deck within 5 blocks of both groups of decision makers (top two graphs); mean proportion of choices from good decks and bad decks of both groups of decision makers (the last two graphs). Each block contains 20 trials.	74
4.5	The learning IGT scores across the five different blocks of the IGT in healthy controls and people living with chronic pain.	74
4.6	Model validation with PPC on the learning scores across five blocks for the ORL model and the VPP model. Model predictions plotted against actual learning scores. Shaded area depicts the 95% HDI of the posterior distribution.	75
4.7	Parameter recovery of the ORL model (the top pannel) and the VPP model (the bottom pannel) for the chronic pain and the healthy participants. Both the means of the posterior distributions of the true and the recovered parameters were being standardized. Dashed and dotted lines reflect 1 and 2 standard deviations in the standardized space, respectively.	76
4.8	Groups-level ORL parameters across healthy controls and people living with chronic pain.	78
4.9	Differences in groups-level distributions for the ORL model between healthy and symptomatic groups. Solid red lines covers the 95% highest posterior density interval (HDI), and dashed red lines marks the 0 point. Values on the left and right sides of each graph are the lower and upper bounds of the 95% HDI of the comparison between the symptomatic and healthy control groups. If 0 point was included in the HDI, we consider there to be a non-significant difference between the two groups. We found a main effect of group on the reward learning rate (A_+), punishment learning rate (A_-), decay rate (K), perseverance weight (β_p), and the difference between the reward and punishment learning rates for the ORL model.	79
4.10	Estimated parameters of the ORL model plotted against four pain measures. Healthy controls are plotted in red and patients with chronic pain are plotted in blue. The r values were calculated between the model parameters and pain measures for the whole sample ($*P < 0.05$, $**P < 0.01$, $***P < 0.001$).	79
5.1	The proportion of trials selecting the option with a lower probability of yielding losses and with smaller loss magnitudes and the proportion of trials shifting when receiving or not receiving losses respectively. Each dot represents a participant and error bars represent 1SEM.	99
5.2	Model validation with PPC on the model-independent measure of the probability of choosing the correct choice for model 1 and model 2. (a) Trial-level model predictions plotted against actual data. Shaded area depicts the 95% HDI of the posterior distribution. (b) Subject-level model predictions compared with actual data, in relation to the identity line. Error bars depicted the 95% HDI of the posterior distribution. (c) Grand average model prediction across trials and participants.	101

5.3 Effects of block type (volatile vs. stable), and relative outcome value (good vs. bad) on learning rate ($n = 38$). It shows the posterior means along with the 95% HDI for the group means (μ) for each learning rate component (i.e. for baseline learning rate and the change in learning rate as a function of each within-subject factor and their two-way interactions). The 95% posterior intervals excluded zero for the effect of block type upon learning rate (i.e. the difference in learning rate for the volatile versus stable task blocks α_{vs}). This was also true for the effect of relative outcome value, that is, whether learning followed a relatively good (no stealing) or relatively bad (being stolen) outcome (α_{gb}). Participants showed higher learning rates during the volatile block than in the stable block and on trials following good versus bad outcomes. 102

5.4 Relationship of model parameters with model-agnostic summaries of behavior. 103

5.5 Effects of the severity of PD and the severity of non-motor symptoms of PD on the cognitive parameters. The posterior means and 95% HDI for the effect of depression (β_d), anxiety (β_a), apathy (β_{ap}), pain (β_p) symptoms, and the severity of PD (β_u) on each of the cognitive parameters were shown respectively. The depressive and anxiety symptoms uniquely and credibly modulated the extent to which learning rate varied between the stable and volatile task block (α_{vs} ; $\beta_d = -0.48$, 95%-HDI= $[-0.88, -0.07]$; $\beta_a = -0.46$, 95%-HDI= $[-0.90, -0.05]$). The depressive, anxiety and apathy symptoms were all significant correlated with the risk preference parameter (γ ; $\beta_d = 0.77$, 95%-HDI= $[0.01, 1.60]$; $\beta_a = 0.88$, 95%-HDI= $[0.09, 1.69]$; $\beta_{ap} = 0.99$, 95%-HDI= $[0.21, 1.81]$). 105

6.1 The workflow of the recruitment process. 115

6.2 The pipeline of the experiment. 116

6.3 The screenshots of the interface of the EMA questions on the smartphone screen. 119

6.4 The screenshot of the task question on the smartphone screen. 120

6.5 (a) Tinnitus symptom variation of one example participant. (b) Correlation between tinnitus loudness and stressfulness. 124

6.6 (a) The proportion of trials choosing the option with a lower probability of yielding losses, (b) the proportion of winning trials, (c) the proportion of choosing the option with smaller magnitude, and (d) the proportion of trials shifting when receiving losses for stable and volatile blocks separately for the two task schedules (schedule 1 = stable block first; schedule 2 = volatile block first). Each dot represents one session of one participant and error bars represent 1 SEM. 125

6.7 Effects of block type (volatile vs. stable), and relative outcome value (good vs. bad) on learning rate ($sessions = 58$). The 95% posterior intervals excluded zero for the effect of outcome value upon learning rate (α_{gb}), i.e. whether learning followed a relatively good (no stealing) or relatively bad outcome (being stolen). Participants showed higher learning rates on trials following good versus bad outcomes. However, they did not improve their learning rates in volatile block versus in stable block. 126

- 6.8 (a) Pairwise correlations between behavioral statistics of task performance across sessions and participants. The percentage of choosing the option with smaller stealing probability (*correct*) and the percentage of winning were strongly correlated *win*. These two metrics were moderately correlated with themselves across sessions. (b) Pairwise correlation between model-derived phenotypes of task performance across sessions and participants. The baseline learning rate α_{base} , the risk preference γ , and the inverse temperature parameter β were strongly correlated with themselves, whereas the subcomponents of the learning rate did not maintain constant over time. 127
- 7.1 Architecture of the RNN model. The model has a GRU layer (the left part), which receives the previous action (the action is one-hot coded), loss and reward x_t and previous hidden states h_{t-1} and a softmax layer (the right part), which outputs the probability of selecting each deck on the next trial. The GRU layer is composed of two gates, i.e. an update gate and a reset gate, containing a set of weight parameters $\{W_r, U_r, b_r, W_z, U_z, b_z, W, U, b_{hw}, b_{hu}\}$ (b_r, b_z, b_{hW} and b_{hU} are not shown for simplicity). The output of the cells of the GRU layer h_t are connected to the softmax layer using black lines. The weight parameters in the softmax layer are represented as V 139
- 7.2 Architectures of the four training paradigms. (a) The most straightforward CL approach is based on centralized data sharing, i.e. requiring that every laboratory share data to a central node and aggregating to train a model; (b) IL approach, in which each laboratory trains the model and passes the learned model parameters to the next laboratory for training on its data, until all have trained the model once. (c) CIL, which involves the repetition of the incremental learning process multiple times. (d) FL, in which all laboratories train the same model at the same time and update their model parameters to a central sever where it is used to create a global model. . . 143
- 7.3 Accuracy of single lab models tested against testing datasets from each of the laboratories. The vertical axis represents models trained on a single laboratory dataset, and the horizontal axis represents the testing dataset of each independent laboratory. AVG is the average accuracy of each laboratory model performance over all laboratories. Overall, the prediction accuracy of each single model on the testing sets, no matter from their own laboratory or other laboratories, are not satisfying. Some of them are lower than the chance probability 25%. 145
- 7.4 The training and validation loss of the ten single models for each laboratory. The training loss decreases over time, whereas the validation loss starts increasing shortly after the first 10 or 20 epochs. 145

7.5 Accuracy of four collaborative models tested against testing datasets from each of the laboratories. The vertical axis represents the collaborative models trained on the training datasets from all laboratories and the horizontal axis represents the testing dataset of each independent laboratory. AVG is the average accuracy of each collaborative model performance over all laboratories. Apart from the random and descending-ordered IL and CL model, the prediction accuracies of the collaborative models were significantly improved compared to that of the single models. 148

7.6 Probability of selecting each deck (average over subjects). SUBJ refers to the data of the experimental subjects; CL-based, FL-based, CIL-based, and IL-based are the on-policy simulations of the four collaborative models trained on the IGT with the same payoff scheme and same number of trials the subjects completed, respectively. CL-based, FL-based and IL-based agents behaved more similarly to human subjects, selecting more good decks than bad decks. 149

7.7 The IGT learning scores across ten different blocks of experimental subjects and simulated model agents. The learning process of human subjects, CL-based and FL-based agents was very similar to each other, overall the learning score progressively increasing over blocks. The IL-based agents were least similar to the human subjects, not presenting obvious learning trend in the process. Although CIL-based agents demonstrated similar trends with human subjects, the magnitudes of their learning score were always lower than that of the subjects. 151

7.8 Probability of switching to different deck based on whether received losses or not in the previous trial, average over subjects. The CL-based agents demonstrated almost the same switching strategy as human subjects. FL-based agents ranked second. The switch probability of IL-based agents was obviously higher than that of the subjects when there was no loss. . . 153

7.9 Off-policy simulations of the two collaborative models. Each panel shows a simulation of 30 trials. The four collaborative models behaved quite similarly in responding to the losses, especially the FL-based and CL-based model. CIL-based model was more sensitive to the losses given by Deck C as deeper dips are caused compared to other models. IL-based model presents the least perseverance with Deck A compared with other models. 154

List of Tables

3.1	Summary of the payoff scheme of the IGT	40
4.1	Demographic information for healthy controls and people living with chronic pain	71
4.2	Models and prior fits	75
5.1	Basic demographic and clinical details of the participants	98
6.1	Basic demographic details of the participants	123
6.2	Leave-One-Out Information Criteria of the Five Models	126
7.1	Parameter settings for the centralized and distributed learning models . .	144

Abstract

From Labs to Real-World: Developing Smartphone-Based Methodologies for Enhanced Phenotyping of Human Decision-Making in Clinical Settings

Lili Zhang

The application of computational approaches in behavioral and cognitive science has advanced our understanding of the biological mechanisms involved in learning and decision-making. However, these advances have not been effectively translated into clinical applications, partly due to the limitations of small-scale, lab-based experiments in capturing the complexity of the brain and its interaction with the environment. The widespread use of smartphones presents an opportunity to overcome the limitations of these lab-based experiments. This thesis aims to develop and validate smartphone-based methodologies that can provide richer data sets and larger samples for studying human decision-making and facilitating clinical translation. Several case studies were conducted to demonstrate the feasibility of these methodologies.

The first step involved moving experimental settings from laboratories to naturalistic settings using lab-in-the-field methodology. Two case studies were conducted: one involving individuals with chronic pain and another involving patients with Parkinson's Disease. These studies revealed correspondences between altered decision-making performance and clinical variables of interest, showcasing the potential of testing decision-making phenotypes outside the laboratory. To capture momentary phenotypes in different contexts and time points, the Ecological Momentary Assessment (EMA) methodology was employed. A proof-of-concept EMA study assessed momentary tinnitus experience and its impact on decision-making using a mobile app, demonstrating the possibility of dense sampling of human decision-making in the daily contexts. Ethical and legal constraints on large-scale human phenotyping were also addressed. The feasibility of training computational models of decision-making using distributed learning strategies was examined using a "many-labs" data pool. Federated Learning offered an alternative when large-scale private data collection using smartphones is restricted.

In summary, this thesis explored smartphone-based methodologies to develop robust neurocognitive models of mental health conditions, implemented through various behavioral and cognitive studies. These methodologies have the potential to complement traditional lab-based experiments, significantly enhancing our understanding of cognitive mechanisms and expediting clinical translation.

Chapter 1

Introduction

1.1 Overview

One of the brain's most fundamental abilities, which is also critical to the success of human beings and other species, is to learn and make adaptive decision-making to exploit the environment for resources and avoid harm [1]. Whether as a professional in a business or in our private personal life, we routinely make series of decisions from choosing which restaurant to have lunch at to deciding how much production capacity installation is required to meet the demand in business, with each decision depending on previous feedback from a potentially changing environment. Decision-making has been investigated by multiple disciplines, from economics [2], to psychology [3], to neuroscience [4], [5], to computer science [6]. The application impact of these studies is of enormous importance since decision-making processes are transversal to multiple and various contexts, including, for example, the medical, political-economic, organizational, and business fields. Particularly, many mental disorders are associated with aberrant learning and decisions. The improvements in understanding the neural underpinnings of decision-making has the potential to revolutionize the diagnosis and treatment for various neurological and psychiatric disorders [7]–[12].

From a cognitive science perspective, decision-making is a cognitive process based on a series of basic cognitive processes like perception, attention, and memory, which involves the accumulation of evidence associated with the utilities of possible options

and then the choice of one of them given the evidence [13]. The most conventional decision-making research in psychology and cognitive science generally begins with developing a theory or hypothesis about what should happen in pre-defined experimental paradigms given to the subjects under highly controlled laboratory conditions, away from realistic settings [14], [15]. The behavioral paradigms are designed to mimic the cognitive operation in the real world. Economic and Reinforcement Learning (RL) theories have been widely applied to formalize the computations taking place in the brain in these simulated scenarios. The computational parameters extracted from such theory-driven models have been utilized as a promising phenotyping tool for humans. Additionally, these parameters can be used to simulate time series of latent cognitive process, which can be linked to neural activities. Data-driven Machine Learning (ML) models have also been utilized to characterise the decision-making process, improve classification of disease and predict treatment outcomes. The continuing and rapid development of these computational approaches holds great promise for understanding the mechanisms of cognitive process and what leads to maladaptive forms of learning and decision-making in clinical population.

While findings in these lab-based studies have significantly advanced our knowledge about how people learn about their environment and how they make decisions about the best course of action, most of them have unfortunately not been translated into major advances in clinical practice. Apart from the fact that behavioral and cognitive research involves both the most complex organ, the brain, and its interaction with a similarly complex environment, a practical problem is that results from traditional small, single sites lab-based studies generally do not permit the generalization beyond laboratory settings to real-world replication attempts [16]–[18]. Conducting large epidemiological-style studies of human behavior and decision-making in the natural world can be full of challenges. There are many extraneous variables that can influence behavior and decision-making and make it difficult to control for all relevant factors in naturalistic settings. For example, human abilities to process information can be easily influenced or interrupted by the contexts they are in as well as their emotional, social, and physical states. This can limit the internal validity of the experiment and make it difficult to draw clear conclusions about

the effects of the independent variables since we can not guarantee that the extent to which an experimental measure accurately and appropriately captures the construct it is intended to measure due to these variables. Furthermore, conducting experiments in real-world settings can be more resource-intensive than conducting lab-based experiments, as it may require more personnel, equipment and devices, and time to set up. Last but not least, collecting data in subjects' daily routines and in large-scale range can raise ethical and legal concerns, such as concerns about informed consent, privacy, and confidentiality.

As a result, the research on how different computational phenotypes manifest in naturalistic settings and change over time and contexts within large samples has been relatively limited in the literature. However, monitoring behavioral changes in daily contexts in real time in large-scale range can provide valuable insights into characteristics of clinical mental disorders that are generally difficult to capture in laboratory settings. It is a necessity to further disentangle complex cognitive processes and to accelerate the translation from basic science findings to clinical treatment. The rapid ubiquity of smartphones makes it possible overcome many of the barriers for measuring behaviors, clinical symptoms and their changes with high momentary resolution in the natural world. In this thesis, we focus on developing and investing new methodologies using smartphones for basic research so that we can shift from traditional small, single site lab-based paradigm to broad, dense, and repeated assessments in the field in large scale. The first step is to move the experimental context from laboratory to the field, i.e. administering the experimental paradigms under unsupervised and uncontrolled environments of people's everyday experience, which is called lab-in-the-field experiments [19]. The second step is to use Ecological Momentary Assessment (EMA) methodology [20] to measure potential associations between momentary psychological states or symptom changes and performance on decision-making in real time on relevant moments. The implementation of these steps mainly rely on smartphones, which allows us to collect data sets remotely with web-based or mobile-based interfaces on subjects' own devices. Smartphones offer a cost-effective and versatile means of scaling up the range of populations investigated in research, thus improving the diversity and representativeness of the sample. By leveraging smartphones, experimental paradigms and

self-report measures can be administered to a large and diverse cohort, simultaneously and across time. In addition to collecting active data, smartphones provide access to a rich source of passive data, including location and accelerometer data, as well as information on email, calendar, microphone, and camera usage. Nevertheless, they have not been involved in this thesis due to various challenges. A persistent issue faced by us is that collecting much larger samples that contain massive amount of personal information is always constrained by strict legal and ethical regulations for data privacy protection. Thus, we turn to the field of secure and privacy-preserving AI approaches for solutions to fill the gap between personal data protection and data utilization for research.

1.2 Research Objectives

The overall objective of this thesis is to facilitate the translation of research findings in the field of behavioral and cognitive science into clinical settings. This is achieved through developing and validating new methodologies that can deliver on substantially richer, multivariate data sets and larger samples for research on phenotyping human decision-making and demonstrate the versatility of using computational models to reveal characteristics of decision-making of various populations in contexts with higher ecological validity but mixed with more noises. This thesis investigated three methodologies leveraging the rapid ubiquity of smartphones and state-of-the-art machine learning techniques. Specifically, it aims to achieve the following three sub-objectives:

- It aims first, to shift the experimental settings of human decision-making studies from laboratories to daily contexts, i.e. assessing subjects' performance on standardized behavioral paradigms in natural states when they are not being observed. This methodology is referred to lab-in-the-field experiments in this thesis, although the definition of this term was not consistent in the literature [19], [21], [22]. People do not always behave the same way when they are in a highly controlled laboratory environment compared to how they behave in their naturalistic settings. The utility of lab-in-the-field in a range of experimental targets can develop knowledge of what is

possible in this area. Importantly, the preservation of the behavioral tasks guarantees experimental control and theoretical tractability. Validating the feasibility of moving conventional lab-based experimental paradigms in the field is the prerequisite to further enhance the depth of assessments, i.e. dense, repeated assessments, which are simply too difficult to be administered using traditional methods in laboratories.

- If it is feasible to move the lab-based paradigms to real-world settings, a natural further step is to move beyond cross-sectional methods toward longitudinal methods so that we are able to make inferences about specific individuals and to monitor longitudinal changes across time within an individual or group during intervention or treatment. There is a lack of within-subject longitudinal assessments of momentary decision-making behavior in contemporary research, which may present the most significant barrier to the generalization of the laboratory findings. Thus, the second objective of this thesis is to capture the momentary decision-making phenotypes in various contexts and time points, as well as the interaction and confounding variables within the same individual using EMA [20], as a new method of data collection designed for assessing and tracking people's thoughts, feelings, or behaviors in real time, is suitable for addressing this issue.
- If it is feasible to repeatedly sample subjects' decision-making and their psychological states in real time in real world settings, the next step is to extend the range of investigated population. Traditional laboratory experiments often involve young and educated samples [23] and the sample size is generally small, which has been criticized for lacking representativeness. Large-scale phenotyping of human decision-making is an important step to test the extent to which our models generalize and account for diversity in populations. Although smartphone-based data collection can greatly reduce the barriers of large-scale research, in practice, there are ethical and legal restrictions on data privacy and protection, which can make it extremely complicated and challenging to collect data that contains sensitive personal information in large-scale range. An alternative way to increase data size and diversity is via laboratory collaboration and data sharing. However, this is not easy either, especially

when involving larger number of laboratories from different legal jurisdictions, due to privacy, ethical, and data regulation barriers [24]. Thus, the third objective of this thesis is to find a solution to overcome the ethical and legal limitations on large-scale human phenotyping.

1.3 Research Questions

The research questions addressed in this thesis arise from the aforementioned objectives. In addition, the thesis adopts a “Bring Your Own Device” (BYOD) paradigm increasingly common in clinical trials with electronic Patient Reported Outcomes (ePRO) requirements [25]. The thesis extends the BYOD approach to human decision-making experimental paradigms along with a methodological flexibility which supports participants responding at their own pace in their daily situation. The popularity of smartphones makes it much easier to implement such a type of experimental design. The method of collecting data with smartphones has yielded significant findings in various fields [26]–[28]. However, the research on the adoption of smartphones to sample computational phenotypes of decision-making is limited.

Two web-based decision-making paradigms that have been widely used for comparing characteristics of decision-making across groups in laboratory experiments were developed. The web-based attribute enables us to deliver the laboratory experimental paradigms via widely accessible URLs in daily environments. In order to allow subjects to report repeatedly on their experiences across real-time contexts, a mobile application called AthenaCX was used. AthenaCX is a platform for rapid design, configuration and deployment of mobile EMA and interventions. It is adopted to develop context-sensitive decision-making experiments, i.e. as the URL to the web-based paradigm can be easily embedded into AthenaCX so that it is convenient to trigger the decision-making paradigms based on the fluctuations of the subjective well-beings or objective context variables (via data from wearable sensors) using AthenaCX. For instance, administer assessments in natural settings at regular intervals or in response to state changes. The development of these interfaces facilitated the implementation of the lab-in-the-field and EMA methodologies. A further

step towards clinical translation is to conduct large-scale phenotyping, which is often hindered by the ethical and legal regulations of large-scale data collection. This problem might be addressed by the rapid development of secure and privacy-preserving distributed training strategies, which makes it possible to fit the model without extracting data from subjects' devices.

Four case studies are conducted in this thesis to prove the feasibility of implementing such methodologies that may facilitate the clinical translation of basic decision-making research. Given the objectives and research context we can now state the research questions of this thesis which are:

- 1. Are we able to preserve the efficacy of the computational phenotypes to distinguish the decision-making strategies of various populations when shifting the research settings from laboratories to real-world contexts?**

This research question explores if it is feasible to administer the web-based behavioral paradigms and cognitive tests without supervision and if the additional noise and potential confounds introduced because of the unsupervised experimental setting undermines the efficiency of the experimental paradigms. This is validated by two lab-in-the-field experiments implemented with smartphones involving various participant populations. Through these case studies, we would like to examine if we are able to collect adequate data sufficiently complete to facilitate modelling (e.g. unpaid participants are more likely to quit tasks that they do not enjoy when they are unsupervised); if the engagement of the participants is good enough; and if we are able to reveal plausible relationships between the variables of interest.

- 2. Is it feasible to repeatedly sample human decision-making at various occasions based on their subjective feelings and psychological states? If it does prove feasible then we can ask if the computational phenotypes of decision-making captured are static traits of an individual or fluctuate with the dynamic momentary symptoms or subjective well-being? Is it possible to identify the causal relationship between changes in symptom fluctuations and decision-making performance?**

Measuring momentary decision-making phenotypes is not easy. There needs to be a trade-off between the length or duration of an assessment and the frequency at which it can be administered. The behavioral paradigms generally involve hundreds of trials, requiring more than ten or twenty minutes of focus to complete. The long duration of testing can cause participants to tire of the assessments, which could decrease the accuracy of responses and compliance, thereby increasing attrition. As a result, researchers have rarely incorporated the behavioral paradigms in longitudinal-based designs, especially when the tasks are administered multiple times in the wild without any supervision. However, humans are constantly influenced by their immediate context and their own physical and mental status, which are very likely to lead to changes in decision-making performance. A cross-sectional study that aims to identify group differences is not capable of detecting the intra-individual variations. We argue that it is worth going beyond cross-sectional designs toward longitudinal designs to understand the variation in the decision-making behavior. Thus, we will conduct a case study to examine the feasibility of such experimental design and try to link the momentary phenotypes to real-time symptom and well-being variables.

- 3. Is it feasible and reliable to fit computational models of human decision-making using the distributed training technique to overcome the barriers for formulating large sample size? Could the distributed training strategy achieve equivalent efficiency in comparison with the traditional centralized training approach?**

This research question aims to examine the feasibility and efficiency of applying distributed learning techniques from the machine learning field to fit models without access to all data. This alternative model fitting approach will be beneficial to overcoming the difficulties which can occur in the process of expanding the size and diversity of the research populations. The answer to this question is particularly interesting given that it is attractive to use passive data generated by the smartphones to create complementary digital phenotypes of decision-making, but it is tricky to collect such data due to the involvement of massive amount of sensitive personal information contained in practice. Distributed learning is a promising method

for overcoming the ethical and data privacy barriers for large-scale phenotyping of human behavior because the final model has access to personal data from all participants. We would like to examine if the distributed learning strategies can achieve comparable performance as the traditional centralized data pooling approach.

1.4 Organization of the Thesis

The organization of this thesis is as follows:

- **Chapter 2** introduces the concepts of lab-in-the-field experiments and EMA methodologies adopted in this thesis. Theoretical models of decision-making are also reviewed in this chapter.
- **Chapter 3** describes the applications implemented for collecting data in more ecologically valid settings and the framework of the computational phenotyping process, which is the main method of data analysis in the following case studies.
- **Chapter 4** involves a case study that implemented the lab-in-the-field experiment methodology aiming to quantify decision-making differences between chronic pain individuals and healthy controls outside the laboratory with a cross-sectional design.
- **Chapter 5** is another lab-in-the-field study. This case study was designed to reveal if the non-motor symptoms of Parkinson's Disease are associated with impaired adaptation of learning to contingency volatility in the decision-making process.
- **Chapter 6** conducts a pilot EMA study aiming to explore if the computational phenotypes of decision-making are traits of an individual or fluctuate with the momentary physical distress, e.g. tinnitus.
- **Chapter 7** examines if it is possible to train a large-scale model of human decision-making without accessing to the raw data directly. Several distributed training approaches designed for privacy preservation were compared to the traditional centralized learning approach.

- **Chapter 8** concludes the thesis. Future directions are also discussed with respect to the different research areas.

Chapter 2

Methodology and Theoretical Foundations

This chapter delineates the definitions of the two methodologies, i.e. lab-in-the-field experiments and ecological momentary assessment, adopted for collecting granular, multivariate, and high-frequency data on behavioral phenotypes and clinical symptoms in real world contexts. The lab-in-the-field methodology emphasizes the importance of examining the efficiency and reliability of the computational phenotypes of decision-making in naturalistic instead of laboratory contexts. The EMA methodology aims to capture the variations of decision-making across time and situations. Given that the implementation of the two methodologies is technologically intensive, we explained the rationale and benefits of implementing the two methodologies on smartphones. We also present a comprehensive review of the economic theories of choice and reinforcement learning models used to build the foundations for understanding the cognitive process of the choice behavior. The recent development on applying Deep Learning (DL) models to understand cognitive process of decision-making is also introduced in the review.

2.1 Methodology 1: Lab-in-the-field Experiments

2.1.1 Decision-making Paradigms

In the past several decades, laboratory experiments have been the most frequently applied experimental pattern in the fields of social, behavioral and brain science that aims to provide an explanation for why human behaviors exhibit systematic patterns [29]. The cognitive process of decision-making, which is fundamental to all human behavior, has also been studied extensively in laboratories. Because laboratory research tends to focus on narrow problems, researchers make a set of assumptions about how people make decisions that are then embedded in the standard paradigms. A diverse repertoire of standardized and validated decision-making paradigms have been developed. Participants in such paradigms are generally required to choose from a set of options in a specific small-scale simulated situation in which the related variables are manipulated. The feature of manipulation of variables enables the researchers to establish cause and effect relationships. Although these artificial paradigms provide an oversimplified model of social situations, they allow for specific adaptation to more closely mirror real-life interactions, for example by having the participants earn their endowment through exerting effort or by having them interact with acquainted others. Importantly, the application of computational models (which will be introduced in the last section of this chapter) to the decision-making paradigms presents a promising new approach for inferring the underlying mechanisms which generate observed behaviors. The parameters derived from computational models constitutes a phenotype that can be used to represent the characteristics of an individual's cognitive mechanisms [30]. Thus, the decision-making paradigms have become effective tools for objectively measuring cognitive processes and the neural basis underlying behaviors in humans and animals, especially in the field of clinical psychology [7], [8], [11].

2.1.2 Laboratory Experiments

Administered in laboratory environments, the factors orthogonal to the theoretical problem being studied in the simulated paradigms such as context and background are all removed

so that the experimenter can ensure appropriate measurement of treatment or manipulation effects and eliminate potential confounds from the study. The laboratory setting offers several important benefits relative to other methodologies when studying human behavior. Since laboratory participants are randomly assigned to different conditions in a controlled setting with little noise, researchers can more accurately measure the effects of the manipulated independent variables on participants' subsequent decisions. This in turn makes it possible to establish cause and effect relationships between variables. The controlled environment also means it is easy to replicate the exact environmental conditions of the original experiment. The artificial environment created in the laboratory is so far from real-life that the results gained from such experiments tell us very little about how respondents would actually act in real life [16]. It is difficult to generalize the conclusions from laboratory settings to complex situations in the real world.

2.1.3 Lab-in-the-field Experiments

Field experiments, which are conducted in naturalistic environments typically without deliberate experimental manipulation of the independent variable and without participants' knowledge, are perhaps the most effective way for achieving the highest ecological validity both in terms of the representation of the stimuli and the research context. However, the gain in the realism of natural experiments is obtained at the expense of the stimuli occurring without restrictions or control of the environment [21], [31], making it hard to replicate the results and make direct comparisons to other environments and populations. As an intermediate approach, lab-in-the-field experiments combines elements of both lab and field experiments, providing researchers with a tool that has the benefits of both, while minimizing the respective costs [19].

Specifically, lab-in-the-field experiments refer to studies conducted in a naturalistic environment targeting the theoretically relevant population but using a standardized, validated experimental paradigm [19]. This experimental method is adopted throughout this thesis to maintain experimental control and theoretical tractability. On the one hand, lab-in-the-field experiments are conducted in naturalistic settings as field experiments, capturing

behavioral features in the real world, which is particularly valuable since people generally demonstrate more complex, heterogeneous behaviors while they are not supervised. On the other hand, unlike field experiments, which are often limited by real-world circumstances thus the results are difficult to be replicated, lab-in-the-field experiments target the relevant population and setting, which increases the applicability of the results. Additionally, the preservation of the standardized paradigm element in the lab-in-the-field experiments permits us to maintain relative tight control, while allowing for direct comparisons across contexts and populations.

2.2 Methodology 2: EMA

EMA is a range of research methods that use monitoring or sampling strategies to repeatedly assess people's emotions, behaviors and experiences at the moment as they occur in natural settings [20]. The collection and assessment of the real-time data about when and under what circumstances an action takes place is of significant importance for revealing the behavioral patterns in real life. The EMA approach overcomes several longstanding problems in behavioral assessment. Firstly, EMA approaches focus on assessing phenomena on a momentary basis. This is often more accurate than conventional methods that require people to recall their behavior or feelings over long periods later, thus causing retrospective biases [32]. The moments can be strategically selected for assessment, whether based on particular features of interest, by random sampling, or by other sampling schemes. The vast majority of the research in decision-making studies is cross-sectional that is only capable of capturing variability between individuals. EMA studies are naturally in a longitudinal design, in which subjects are required to complete multiple assessments, providing a picture of how their experiences and behavior vary over time and across situations. The repeated measurements represent the dynamic nature of the psychological processes of a single individual, which are essential for developing and testing causal models of human behavior.

The EMA methodology is overlapped with the lab-in-the-field experiments methodology to some extent. The lab-in-the-field methodology is mainly used to move the

experimental context outside the lab and they generally indicates a cross-sectional design in this thesis. EMA methodology places emphasis on sampling subjects' behavior and feelings repeatedly at relevant moments in a longitudinal design, allowing us to test possible links between brain and changes in symptoms over time. Until recently the measurements in EMA were mainly restricted to self-reported data, however, their application via smartphones has broadened the assessments to include data collected from their integrated sensors or paired wearables. Decision-making paradigms seldom occur in EMA-based studies to the best knowledge of the author, but they are incorporated in the EMA study of this thesis to detect the fluctuation of the cognitive functioning. To this end, the essence of EMA methodology in this thesis is in fact an extension of lab-in-the-field experiments, i.e. administering the lab-in-the-field experiments multiple times according to the variations psychological status and contexts. This kind of experimental design enables us to capture the potential sources of within-person variation in terms of the cognitive processes across time.

2.3 Methodology Implementation

Both methodologies, especially the EMA methodology are technologically intensive. There are many ways to implement the behavioral paradigms and sampling schemes. The exponentially increasing use of smartphones presents a promising tool for moving the experiments beyond laboratories towards peoples everyday lives and achieving complex structures of sampling schemes. They are naturally chosen as the devices where we implement the lab-in-the-field and EMA experiments. Since the EMA experiment is an extension of the lab-in-the-field experiment, it is more complex and difficult to conduct. Therefore, we here mainly discuss the key advantages of implementing EMA methodology using smartphones.

2.3.1 Historical Methods

The implementation of EMA studies, i.e. how to engage participants during their daily and private lives "without disrupting the phenomena to be observed", is one of the most complex challenges faced in EMA studies. Historically, there have been two main approaches for implementing EMA designs. One approach involves a signalling device of some sort and paper-and-pencil surveys. The second approach combines signalling and item presentation using a palmtop computer [33]. The former method has been achieved using a variety of tools. Initial research required participants to carry pagers [34] that would beep periodically throughout the day. Participants then would complete paper-and-pencil surveys that they also carried out [35]. The desire to reduce participant burden led to some researchers using digital wristwatches with preprogrammed alarms that could alert participants to complete a survey card [36]. Nevertheless, carrying packs of survey cards instead of paper diaries is still a heavy burden for the subjects, thus easily undermining their compliance. The invention of the Personal Data Assistant (PDA), which was a small personal computer, addressed this issue because it could be preprogrammed both to signal participants and provide the survey interface. PDAs also offered greater control over the data collection process, such as the ability to limit access to the PDA between responses and time stamp responses [37]. However, the problem with PDA approaches is that it requires researchers to purchase the hardware upfront, give it to participants, and recover it after data collection.

2.3.2 Smartphones

Over the past decade, as technology has advanced, PDAs were largely replaced by smartphones and tablets that have opened a whole new world for EMA studies [26], [27], [38]. This is the BYOD phenomenon referred to earlier and common in an ePRO context. Firstly, this approach drastically lowers researchers' hardware costs because participants can use their own devices. Secondly, smartphones are sensor-rich, computationally powerful, and nearly constant companions to their owners, making them capable of capturing numerous interacting and confounding passive variables within the same individual. This is an unparalleled superiority compared with other methods. Furthermore, it is easy to query

people about their subjective feelings, thoughts, and behaviors via the notification system of a smartphone. Last but not least, smartphones are compatible with both mobile and web applications, both of which are suitable for implementing the decision-making paradigms that are crucial elements both in lab-in-the-field and EMA methodology in this thesis.

There are multiple ways in which smartphone-based assessment can be employed in EMA studies. One popular approach uses the Short Message Service (SMS) on smartphones (e.g. SurveySignal [39]). The SMS method uses text messages as a signaling device by texting the participant a hyperlink to an online survey, which is accessed through their smartphone's web browser [40]. Text messages have the advantage that they are relatively easy and inexpensive to deliver and are not limited by the model of an individual's phone [41]. Individuals will likely be familiar with sending and receiving text messages and researchers only need to know how to create an online questionnaire using their favorite survey program (e.g., Qualtrics, Survey Monkey [39]). One of the downsides of this approach, however, is that participants must have cellular connectivity to receive signals and respond to new surveys. If participants are in an area with no cellphone service (e.g., the subway, the basement of a building), they will receive their signals much later than intended or may even miss the survey altogether [42]. Additionally, this approach is only limited to self-report assessments, which means it is not suitable for researchers wishing to capture other sources of data via onboard and external sensors or to integrate the assessment with user activities.

Another approach using smartphones is to collect data with an application installed locally on participants' smartphones, which allows for greater flexibility and ease of use and thus has become increasingly widespread in recent EMA literature [40], [43]–[46]. Similar to the PDA approach, the device itself is used for both signalling and data collection. Depending upon how the app is programmed, it may not require constant cellular connectivity. This approach also allows researchers to easily access to other device data, e.g. microphone, location, accelerometer etc. so that more related passive data of the same individuals can be collected. In work related to this thesis, we had access to a no-code platform for rapid design and deployment of complex EMA studies.

2.3.3 Data Quality

Research on smartphone data collection is still in its infancy, we discuss internet-based experiments more generally in terms of data quality in this section. A major concern about internet-based experiments is whether data collected online yield trustworthy and equivalent results compared to those generated in lab-based settings. Web-based data are thought to be noisier than in-person data. Such noise could result from numerous sources, including not only lack of diligence from participants, but also the influence of a variety of situational and personal confounding variables in the field, which are generally removed in the lab. For instance, it is natural to worry about whether online participants are more careless about their performance because of a lack of incentives; whether the anonymity provided by internet participation could result in deceptive responses; whether the remote participants lack focus while they are unsupervised as in traditional laboratory settings.

Reassuringly, little evidence has turned up in support of these concerns, either. On the contrary, it was suggested in comparative studies that internet participants appear to be no less motivated than their traditional peers [47] and the cloak of internet anonymity is more of a benefit than a detriment [48] in terms of answering online surveys. Further studies have shown that data from web-based samples need not reduce data quality, even for performance-based, and stimulus-controlled experiments [49]. Apart from more traditional web-based methods, several studies have been conducted to demonstrate the reliability and validation of internet-based cognitive testing via crowdsourcing platforms, such as Amazon Mechanical Turk [50], [51] and Prolific [52]. Amazon Mechanical Turk was designed for individuals and businesses to outsource their processes and jobs to a distributed workforce who can perform these tasks virtually, while Prolific was explicitly designed for the scientific community. Such online labor markets have become enormously popular as a source of research participants [53], [54].

The utilization of smartphones for gathering behavioral choices could generate even noisier data because the testing environment is less controlled. However, it turns out that the mass collection of experimental data using smartphones could outweigh the inherent concerns regarding collecting data outside a controlled laboratory setting (see [55] for

a comprehensive review). Smartphone data showed similar effects to those observed in the lab and with comparable quality to in-person studies across multiple domains of cognition [56]. For example, using computational modelling analysis, the Great Brain Experiment app replicated the lab-based finding that self-reported affective state is not explained by current earnings, but by the expectations and prediction errors arising from those expectations [57].

2.4 Economic Theories of Decision-making

Decision-making is a primary target of investigation in the field of computational modelling. Various theories have developed in the field of economics to explain how people evaluate the utilities and values of options presented to them, while reinforcement learning (RL) [58] theories in computer science provide valuable frameworks to model how decision-making strategies are tuned by experience. Additionally, researchers from neuroscience have begun to elucidate the mechanisms of neural substrates in the brain responsible for the computational steps of learning and decision-making [4], [5], [59]. Not surprisingly, economic and RL theories are now frequently featured in the neuroscience literature, and play an essential role in contemporary research on the neural basis of decision-making. Thus, computational models of decision-making not only provide a formalized description of how decisions are made, but also presents more general possibilities for delineating the neural underpinnings of decision-making.

2.4.1 Economic Theories of Choice

Decisions under Uncertainty

Uncertainty is pervasive in many situations we encounter in daily life. From small choices to large lifetime decisions, we rarely know for sure the consequences that will stem from our choices.

Knowledge about the likelihood of possible outcomes of a choice can lie anywhere on a continuum space, from complete ignorance at one end, to various degrees of ambiguity,

to risk, to certainty at the other end [60]. Ambiguity refers to a condition where the probabilities of different outcomes are not available, whereas in a risk condition, the outcomes of choice options and the outcome probabilities are both known to the decision-makers. Here, we first focus on the economic assumptions on how people deals with the uncertainties in risk conditions while making decisions. In a standard risk paradigm, a decision maker is presented with an explicit description of probability distributions over outcomes. There are two key components in the economic models: Stimulus Values (SV) computation and a comparator process. SV is a measure of the expected benefit of consuming the different options. It is assumed that decision-makers will assign a value to available choice options in order to compare the options and select the one that is most advantageous. One natural option of valuation is to decompose choice options into possible outcomes (or rewards) and their associated likelihoods of occurrence. The assumptions to the comparator process ranges from normative models which provide accounts of what should be chosen, e.g. maximizing expected value or subjective expected utility, to descriptive models that describe what is actually chosen, e.g. prospect theory.

2.4.2 Expected Value Theory

For simplicity, we take the simplest gamble where a set of outcomes and their associated probabilities are provided as an example. Imagine simply choosing between a sure payment of \$49 and a 50-50 lottery of receiving \$100 or nothing. The expected value of the risky (or gamble) stimuli is the same for all decision-makers, as it uses the objective outcome amount and probability level, i.e. multiplying each outcome by its probability of occurrence and adding overall probability-outcome pairs:

$$EV(X) = \sum_x p(x) \cdot x \quad (2.1)$$

where x is the payoff for each outcome, $1, \dots, X$, and p is the probability associated with that outcome. The expected value theory advice about how to choose between two gambles was simply to choose the one with the highest expected payoff. Thus, in our case, decision makers who apply this rule will prefer the gamble ($\$50 = 0.5 \times 100 + 0.5 \times 0$) to a sure

payment of \$49 because the expected value of the gamble is higher than the value of the sure thing.

However appealing the idea of weighting payoffs by their probability might have been in this early expected value theory, expected value maximization is not a universally applicable decision criterion because it could not explain the risk aversion behavior, e.g. why anyone would purchase insurance. A risk-neutral decision maker is indifferent between a gamble and its expected value; a risk-seeking decision maker pursues a risky prospect instead of a sure payment. On the contrary, a risk-averse person prefers a sure payment over a risky prospect of equal or high expected value. However, the expected value model does not allow decision makers to exhibit risk aversion behavior, i.e., no person would prefer a sure \$49 over a lottery with a 50% chance of winning \$100 or nothing under this assumption, but the fact is it happens to many people.

2.4.3 Expected Utility Theory

In order to deal with the problems of the expected value theory, Bernoulli proposed that instead of evaluating options by their objective payoffs, people utilize their expected utility or “moral value” to make decisions. Thus, in Bernoulli’s model, decision makers choose the option with the highest expected utility:

$$EU(X) = \sum_x p(x)u(x) \quad (2.2)$$

in which $u(x)$ is a function that maps the objective payoff x , on the x-axis into the utility of obtaining outcome x and this function is no longer linear but usually "concave", for example, a power function of the form $u(x) = x^\theta$. The exponent element θ in this function is a free parameter that varies between individuals. It is used to describe the curvature of the function and importantly represent an individual’s degree of risk aversion. Because of the introduction of this parameter, the expected utility of a risky choice option can exhibit differences between decision makers since the same objective outcomes can map into different levels of subjective utility. Individuals with smaller θ value ($\theta < 1$) are more cautious when resolving choices among options that differ in risk, while those with higher

θ value ($\theta > 1$) are willing to take on great risks in the hope of even greater returns. As a result, a risk aversion decision maker ($\theta < 1$) believes the utility gained by receiving \$50 is more than half the utility gained by receiving \$100, thus prefers \$50 for sure to a 0.5 probability of winning \$100.

This idea of expected utility became a central assumption in the economic analysis of choice under risk and uncertainty since it was proposed. However, it did not last a long time before it was questioned. One of the most celebrated challenges was known as the Allais paradox, which was initially proposed in 1953. In 1979 [61], the Allais paradox was formally incorporated into the study of behavioral economics by Amos Tversky and Daniel Kahneman, who presented a simplified version of the puzzle in their book:

Decision 1: choose between (A) a 61% chance to win \$520,000; (B) a 63% chance to win \$500,000.

Decision 2: choose between (C) 98% chance to win \$520,000; (D) a 100% chance to win \$500,000.

The majority of people chose option A over B in the first decision and option D over C in the second decision, which violates the utility maximization principle. The violation derives from the willingness of taking an additional 2% risk in exchange for an extra \$20,000 in potential winning for one decision but not for the other. In other words, people were risk-seeking and chose the option with a greater expected value for the first decision but were risk-averse and chose the option with a lower expected value for the second decision, despite the possible outcomes being of equal value for both problems. People explained their choice in **Decision 1** as a preference for a higher possible prize given that the difference between the probability of 61% and 63% is very small. They explained their choice in **Decision 2** as a preference for a certain prize over a slightly smaller prize that entails a small possibility of receiving nothing. Apart from the Allais paradox, there were a significant number of observations that could not be accounted for by expected utility. Tversky and Kahneman [62] revealed a four-fold pattern of risk attitudes: risk seeking for low-probability gains and high probability losses, coupled with risk aversion

for high-probability gains and low-probability losses. Choices consistent with this fourfold pattern, but violate expected utility maximization were observed in several studies [63], [64].

2.4.4 Prospect Theory

While there had been various attempts to account for each of the failures of expected utility, prospect theory [64], [65] accommodated neatly most of the known failures. It has become the leading behavioral model of decision-making under risk and the major work for which Daniel Kahneman was awarded the 2002 Nobel Prize in economics. Kahneman [64] improved the expected utility theory mainly from three aspects: (1) introducing a transformation mapping objective probabilities to subjective probabilities; (2) defining subjective values not on total wealth as in expected utility but rather on gains and losses relative to a dynamic reference point; and (3) allowing losses to have a different mapping into value than that of gains, a phenomenon they called loss aversion. According to prospect theory, the value V of a simple prospect that pays $\$x$ with probability p (and nothing otherwise) is calculated by:

$$V(x, p) = w(p)v(x) \tag{2.3}$$

where v represents the subjective value of the objective outcome x . $w(x)$ is a weighting function that measures the influence of the probability p on the attractiveness of the prospect. Prospect theory replaces the utility function $u(\cdot)$ over states of wealth with a value function $v(\cdot)$ over gains and losses relative to a reference point, with $v(0) = 0$. This famous value function in the 'S' curve, i.e. concave for gains and convex for losses and steeper for losses than for gains, addresses the flaws in expected utility theory. The most famous value function proposed by Tversky and Kahneman [62] (Figure 2.1) relies on a power function:




Figure 2.1: A representative value function proposed by Tversky and Kahneman [62].

$$v(x) = \begin{cases} x^\alpha, & x \geq 0 \\ -\lambda(-x)^\beta, & x < 0 \end{cases} \quad (2.4)$$

where $\alpha, \beta > 0$ are two free parameters that govern the curvature of the value function for gains and losses, respectively, and λ is a loss aversion parameter that determines the sensitivity of losses compared with gains. Thus, the value function for gains (losses) is increasingly concave (convex) for people with smaller values of α (β) < 1 , and an individual with a larger value of λ is more sensitive to losses than gains. These parameters constitute a computational phenotype [30], which could vary within and between individuals, providing an opportunity to examine individual differences in computational mechanisms. With this particular set of parameters, different risk attitudes people have for gains and losses could be explained. Additionally, the reference point implies that people tend to focus less on their final income or wealth, and more on the relative gains or losses that they will get.

The major difference between the decision weights w and objective probabilities is observed with extreme probabilities (either very low, e.g., 1% or very high, e.g., 99%). In general, people tend to overweight low-probability events in judgment, which helps to explain the irrational appeal of gambling and insurance for very low-probability events. On the other extreme, people tend to be underweight highly certain outcomes. People will pay much less, for instance, for a lottery in which they have a 99% chance of winning \$1,000 than they will for a lottery in which they have a 100% chance of winning \$1,000,

but there is little difference between the amount people would pay for a 50% versus 51% chance of winning \$1,000. The objective difference in probabilities (1%) is identical, but its impact on one's decision is not. The most popular form of the weighting function [66] assumes that the relation between w and p is linear in a log-odds matrix:

$$w(p) = \frac{\delta p^\gamma}{\delta p^\gamma + (1-p)^\gamma} \quad (2.5)$$

where $\delta > 0$ represents the elevation of the weighting function and $\gamma > 0$ measures its degree of curvature. The weighting function is more elevated (exhibiting less overall risk aversion for gains, and more overall risk aversion for losses) as δ increases and more curved (exhibiting more rapidly diminishing sensitivity to probabilities around the boundaries of 0 and 1) as $\gamma < 1$ decreases.

2.4.5 Neural Representations of Stimulus Values

Although the economic models of decision-making were not designed to reveal the underlying processes implemented in the brain, cognitive neuroscience literature has assumed that the hypothesized SVs are encoded either in single neurons or in populations of neurons within a brain region. Under this assumption, the firing rate of such units in every trial of choice, or the activity level of such regions, should be proportional to the subject- and stimulus-specific SV measures that have been reviewed above. It has been proved in many studies that SVs are represented in a specific set of brain regions. Although theoretically this prediction can be tested using single unit neurophysiology, blood-oxygenation-level-dependent (BOLD) signal from functional Magnetic Resonance Imaging (fMRI), electroencephalography (EEG), or magnetoencephalography (MEG) to look for neurons or brain regions in which the measures of neural activity correlate with the inferred SVs, it is not an easy problem that requires much thoughtful experimental design and numerous controls. Despite the challenges, the body of data available today, taken as a whole, provides all of the necessary checks and has led to a robust set of findings regarding the implementation of computing SVs in the brain.

Instead of inferring the SVs with parametric functions in expected utility theory or

prospect theory, initial studies simply ask subjects to provide a value for each choice, either using liking ratings or incentive-compatible bids. In [67], hungry subjects were shown a picture of a familiar food snack on every trial and had to decide how much to bid for the right to eat that snack food at the end of the experiment. The bids provide a behavioral measure of the SVs on every trial. The study revealed that the neural responses as measured by BOLD fMRI in the ventromedial prefrontal cortex (vmPFC) and right dorsal lateral prefrontal cortex correlated with the bids. In [68], subjects were shown 50/50 gambles involving both a potential monetary gain and a potential monetary loss and were asked to choose between them and a fixed reference payoff of \$0. It was found that a similar area of vmPFC and ventral striatum correlated with the value of the potential gains and losses. Similar results have been found in dozens of follow-up studies [69]–[72]. Together, these studies provide convergent evidence for the hypothesis that the vmPFC is involved in the computation of SV signals during choice.

As far as uncertainty and risk influence computation of SVs and SVs are represented in a set of brain areas, uncertainty- and risk-mediated changes in value would be expected to be represented in these brain regions, which is indeed the case [72]. Several regions have been found to encode the two main components of outcome-probability decomposition of risky options. Some findings in animal experiments [73], [74] suggested that the phasic responses of dopamine neurons were indicative of value coding reminiscent of expected value and expected utility. Accordingly, across cells, the sensitivity to probability correlates with sensitivity to magnitude. However, some cells were more sensitive to one than the other component. The primary target regions of dopamine neurons are the striatum and prefrontal cortex. Functional neuroimaging has revealed that value-related BOLD activations in the striatum were well-captured by a value function of the prospect theory [68] and show features not contained in the many alternative theories.

The prefrontal cortex is connected to both the striatum and the dopaminergic midbrain. Simultaneous recordings from neurons in the lateral prefrontal cortex (IPFC), orbitofrontal cortex (OFC) and anterior cingulate cortex (ACC) have revealed reward probability and magnitude signals in all three regions [75], [76]. The coding of both outcome and probabil-

ities in single ACC, OFC and IPFC neurons could form the neural substrate of risky option value computations such as the expected value or expected utility combination of outcomes and probabilities. Under this assumption, neurons or regions closer to the motor output are also affected by the combined coding of probability and magnitude. One example comes from the supplementary eye field (SEF) [77], a region involved in processing eye movements that are also innervated by the OFC. SEF neurons respond in an increasing or decreasing fashion to increases in probability and magnitude during saccade preparation. In summary, progresses in cognitive neuroscience has not only enhanced our confidence in the validation and correctness of the theoretical assumptions of stimulus valuation and choice processes but also facilitated the development of these theories.

2.5 Reinforcement Learning Theories

2.5.1 Decisions from Experience

It has been increasingly appreciated that learning plays an important role in decision-making. Decision-making in the real world involves repeated rather than one-shot decisions and people need to estimate the possible outcomes and probabilities based on learning from feedback rather than a written description as used in the decision-making under risk paradigms. Thus, an alternative experimental paradigm developed is based on decision-making via experience. Under such a paradigm, participants learn about the different outcomes of available choice alternatives and probabilities by repeatedly sampling them and experiencing their consequences. Positive consequences increase the likelihood that the option is chosen again, whereas negative consequences decrease it. The paradigms of understanding decision-making through experience enable us explain more complicated scenarios of decision-making in real life, which also raised the needs of new theories and models to explain the incremented elements involved in such decision paradigms.

2.5.2 Rescorla-Wager Model

Consider a scenario where a gambler standing in front of a row of K slot machines and trying to conceive a strategy for which machine to play, for how many times, and when to switch machines in order to increase the chances of making a profit. The rewards specific to each bandit are dispensed according to a probability and is initially unknown to the gambler. Computer scientists abstracted this problem as reinforcement learning which concerns how an agent, such as the gambler, can learn by trial and error to make decisions (choose slot machine in this case) in order to obtain maximum rewards. A series of algorithms have been developed to accomplish this goal in computer science. In the past few decades, such computational theories are also increasingly becoming central in the field of neuroscience and psychology where the researchers have identified the neural substrates for similar learning and decision functions in humans and animals [78]–[80].

The striking correspondence between RL algorithms and evidence from neuroscience and psychology about how the brain solves the RL problem are not surprising because the development of RL drew inspiration from psychological learning theories. One of the most celebrated computational model of RL derived from a family of phenomena observed in classical conditioning in psychology. This phenomenon only involves learning, but not decisions. However, it is an important subcomponent of the full RL problem, because choosing between actions can be based on predicting how much reward they will produce. In the famous experiment on the salivating dog, Pavlov [81] exposed dogs to repeated pairings whereby an initially neutral stimulus, such as a bell, accompanied food, such as meat powder. He observed that, following the training, the dog would salivate in response to the sound of the bell in the same way it did to the food. Connecting new stimuli to innate reflexes in this way is now called classical, or Pavlovian conditioning, which is one of the core empirical regularities around which psychological theories of learning have been built. Pavlov [81] hypothesized that this conditioned response emerges because a preexisting anatomical connection between the sight of food and activation of the salivary glands comes to be connected to bell-detecting neurons by experience. This very general idea was first mathematically formalized when Bush and Mosteller [82] and extended by

Rescorla et al. [83]. They proposed that learning in classical conditioning was based on a prediction error, the discrepancy between what reward the organism experiences on a particular trial r_t , and what it had expected on the basis of its previous learning V_t :

$$PE_t = r_t - V_t \quad (2.6)$$

If the experienced reward is different from its expectation, a prediction error exists, and the organism should update the expectations of future reward V_{t+1} . By contrast, if the experienced reward is exactly the same as expected, there is no prediction error, the organism should keep their expectation unchanged. Specifically, the expectation of future reward is a function of current expectations (V_t) and this prediction error PE_t , multiplied by a subject-specific learning rate (α) that modulates the influence of the prediction error on learning:

$$V_{t+1} = V_t + \alpha * PE_t \quad (2.7)$$

This equation is known as Rescorla-Wagner (RW) model [83], in which an organism is assumed to maintain a set of predictions of the reward associated with each stimulus s , and these predictions determine the organism's conditioned response to whichever stimulus is observed. The value of stimuli that are not observed remains the same, i.e. $V_{t+1}(s) = V_t(s)$, for all $s \neq s_t$. In essence, what this error-driven model does is compute a weighted running average of previous rewards across previous trials. In this average, the most recent rewards have the greatest impact, whereas the impact of prior rewards declined exponentially in their lag. The learning rate parameter can be equivalently seen as controlling the steepness of the decay, with higher learning rates producing averages more sharply weighted toward the most recent rewards [84] (Figure 2.2).

The RW model of classical conditioning is critically important because it is the first use of iterative error-based rule for the computational modelling of learning. Additionally, it forms the basis of all modern approaches to study how animals learn stimulus-related relationships when the experience with that stimulus is accompanied by other stimuli. It successfully explained the blocking effect by specifying that when multiple stimuli




Figure 2.2: The influence of recent outcomes on choice for three different learning rates ([84]). Red represents a participant with a relatively high learning rate of 0.9. There is a strong influence of the outcome on the recent previous trial ($n - 1$), and a weaker influence on the outcome received two trials back ($n - 2$) but virtually no influence on earlier outcomes. A learner with smaller learning rates, here shown for 0.5 (grey) or 0.2 (blue), shows increasingly longer-lasting influences of outcomes received on trials further back from the current trial.

are observed, the organism makes a single net prediction that is the sum of all of their predictions. Since then, a broadening field of computational modelling has extended the family of RW models to reinforcement learning, tackling some of its limitations (such as capturing second-order conditioning), and still grappling with others (such as reinstatement).

2.5.3 Transforming Values into Choice

After the SVs are determined, there needs to be an algorithm for determining which option to select. The most straightforward solution provided by early expected utility and prospect theory was to select whichever choice has the highest subjective value. However simple this solution might be, it is not consistent with what is observed when people select between options. A well known observation is that people do not seem to always choose the better of two options, even when accounting for differences in the subjective weighting of each of the chosen attributes. If they did, we would expect to see that their choices followed a step function, as long as one option has a higher value than the other, they always choose




Figure 2.3: The impacts of inverse temperature on the steepness of the softmax function.

the former one and vice versa. Instead, people's choices follow a sigmoid-like pattern, more step-like when there is a large difference between the option values, but, as that difference narrows, people start to choose the higher-valued one with less consistency (i.e., are increasingly likely to choose the "objectively" lower-valued option).

To account for these inconsistencies, choice models incorporate a noise term, which determines the extent to which choices will be strictly determined by the higher-valued option versus by random chance:

$$P(C = A) = \frac{1}{1 + e^{-\beta(V_A - V_B)}} \quad (2.8)$$

This function, referred to softmax function, generates the probability of choosing each option. The results are increasingly likely to favor the high-reward option when the difference between the options are greater and increasingly likely to choose either option equally when the difference is small. The steepness of the transition is determined by the inverse temperature parameter β (see Figure 2.3 about how the inverse temperature

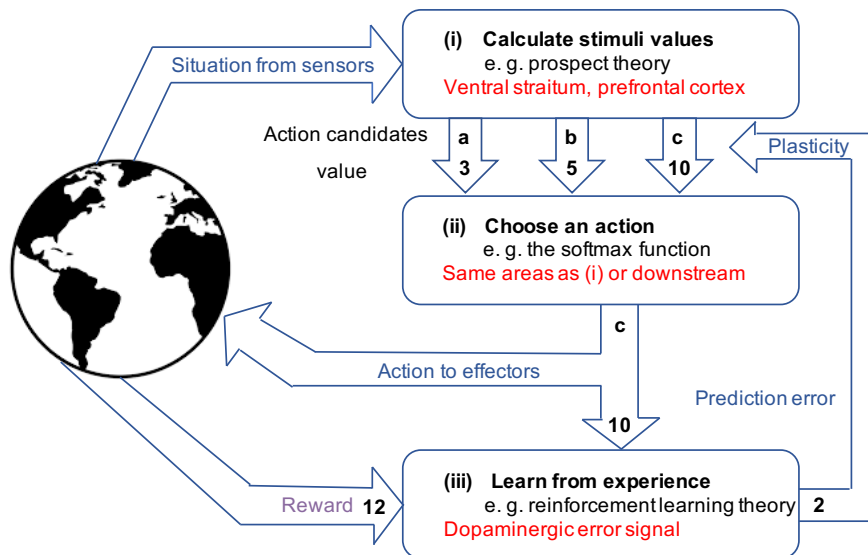


Figure 2.4: The three stages of the theoretical accounts of experience-based decision-making process proposed in [80]. (i) Calculate the values of the available stimuli (here a, b, c) in the current situation. (ii) Choose and execute one by comparing the expected rewards of each option. (iii) Finally, learn from the reward prediction error to improve future decisions. Numbers indicate the predicted action values, the obtained reward, and the resulting prediction error.

controls the steepness of the softmax function). This function can also be generalized to choices between any number of options and is regularly used to model choice in learning models described above.

In summary, the theoretical accounts of the learning and decision-making process can be abstracted in three stages [80]: first, calculate the values of the available stimuli using for example prospect theory; second, choose an action that maximizes the predicted expected value, generally taking the choice noise into consideration by applying a sigmoid function; third, learn from experience to update the expected values of the stimuli. Figure 2.4 demonstrates the three stages and the models and the neural substrates that support these functions.

2.5.4 Responses of Dopamine Neurons to Prediction Error Signal

In the earliest study of reward processing in the brain, Olds and Milner [85] found that rats would continue to press a lever in return for nothing more than a brief pulse of electrical stimulation in a particular region of the rat brain. Since electrical stimulation in the brain area was all that was needed to reinforce the lever-pressing behavior, which is a typical

manifestation of the function of rewards in learning and approach behavior, and since the rate of reinforcement was comparable to that produced by natural rewards, such as foods and water, the authors inferred that stimulation in this brain area somehow induces learning and approach behavior, thus linking brain function in a causal way to behavior. Subsequent studies [86] confirmed that about half of the effective locations for electrical self-stimulation were connected with dopamine neurons that are located in the midbrain and extend their axons several millimetres into the striatum, frontal cortex, amygdala, and several other brain regions. Additionally, dopamine neurons could generate strong phasic activation when encountering rewards, without being stimulated by electric currents. In one of the examinations of dopamine neurons in behaving primates, a strong phasic activation occurred in the dopamine neurons only when the animals found food hidden within a box. The higher the reward, the stronger the dopamine response [87].

However, a closer look revealed that the dopamine neurons not only responded to the reward itself, but also to prediction error [74]. In other words, dopamine neurons resembled the prediction error signal proposed in the RW model, the firing rate of dopamine neurons increasing with the degree to which the reward was expected. In [74], the animals were trained to learn the task in which five different visual conditioned stimuli predicted delivery of liquid reward with different probabilities, ranging in steps of 0.25 from certain delivery ($p = 1$) to certain non-delivery ($p = 0$). Surprisingly, phasic dopamine responses to a reward delivery were zero for the always rewarded stimulus, one for the never-rewarded stimulus, and something in between for the others (see Figure 2.5). Moreover, when rewards were not delivered, dopamine neurons showed a phasic decrease in firing at the time reward would have been expected. Taken together, the activation of dopamine neurons at the time of reward or non-reward was well explained by the RW prediction error. This hypothesis was further tested by Bayer and Glimcher [88]. Recall that the prediction error of the RW model is the subtraction between the current reward and the weighted average over previous rewards, with the weights exponentially declining over trials. To quantitatively examine to what extent the reward history predicts the activity of the dopamine neurons, Bayer and Glimcher used a task in which reward predictions shifted




Figure 2.5: Rasters and histograms from a primate dopaminergic neuron in a classical conditional experiment in [74], illustrating the responses to the conditioned stimuli and reward at various reward probabilities. The five panels correspond to five cues trained with stochastic reinforcement at different probabilities, increasing from top to bottom. Dopamine neurons responses increase with the size of the prediction error.

over trials in conjunction with regression analysis to estimate the weights that best explained the firing rates of the neurons at the time of the reward. The weights estimated suggested that activities of the dopamine neurons reflected the iteratively computed prediction error between the weighted average of the previous rewards and the magnitude of the current reward of the RW model, i.e. they were positive for the current reward, and negative for the preceding rewards, decreasing over trials with a roughly exponential shape. Data from human fMRI experiments also corroborated the notion that dopamine neurons encode a prediction error, in which the BOLD responses in the striatum area resembled the same features of the prediction-error signals seen in animal dopamine recordings [89], [90].

2.6 Deep Learning Models to Predict Human Behavior

Apart from the aforementioned theory-driven models, a more flexible alternative model that can characterize human decision-making achieved through DL models. The unique feature of this type of model is that they do not require tweaking and manual engineering like the theory-driven models [91], [92]. They are being applied in the last case study of this thesis to test the efficiency of distributed learning techniques for privacy perseverance. Dezfouli et al. [92] proposed a RNN model to the problem of human decision-making in the two-armed bandit task (see Figure 2.6 for the structure of the proposed RNN model). The model is composed of an LSTM layer (Long short-term memory, [93]) and a softmax layer. The inputs to the model on each trial are the previous action and the reward received after taking the action, and the outputs of the model are the probabilities of selecting each action on the next trial. The LSTM layer is used to track relevant information regarding the history of past rewards and actions and the softmax layer is used to turn this information into subsequent actions through the output of the softmax function. There are two nodes in the softmax layer since there are two options in the task. The way in which the network learns and maps experiences to future actions is modulated by weight parameters in the network. Thus, the aim of training such a network is to tune the weights so that the network can predict the next action taken by the subjects, given that the inputs to the network are the same as those that the subjects received on the task.

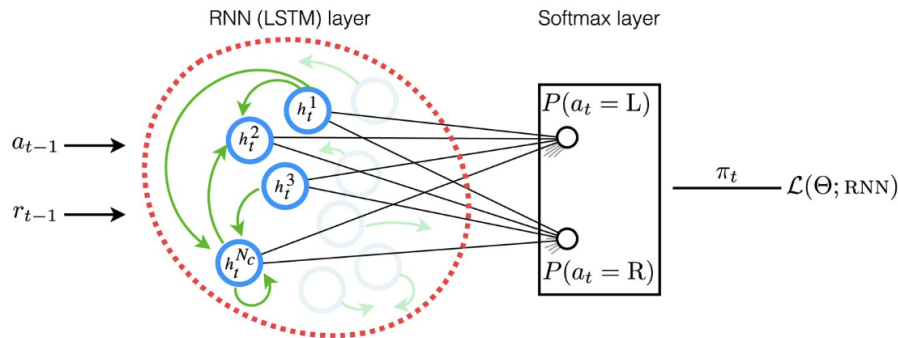


Figure 2.6: Structure of the RNN model proposed in [92]. The LSTM layer circled in red dashed line receives the previous action and reward as inputs, and is connected to a softmax layer (inside the black rectangle) which outputs the probability of selecting each action on the next trial. The LSTM layer is composed of a set of LSTM cells (shown by blue circles), which are connected to each other (shown by green arrows). The output of the cells (denoted by h_t^i for cell i at time t) are connected to a softmax layer using a set of connections shown by black lines. The free parameters of the model are denoted by Θ , and $(\Theta; RNN)$ is a metric which represents how well the model fits subjects' data and is used to adjust the parameters of the model using the MLE as the network learns how humans learn.

However, unlike the theory-driven models, the RNN model has a much larger number of parameters that are hard to interpret. This renders RNNs impractical for studying and modelling individual differences. Therefore, Dezfouli et al [94] developed a new model which is capable of representing individual differences in a low-dimensional and interpretable space, while maintaining the flexibility of RNNs (see Figure 2.7 for the structure of the model). To achieve this, they proposed an end-to-end learning framework in which an encoder is trained to map the behavior of subjects into a low-dimensional latent space, which quantifies aspects of individual differences along different latent dimensions (2.7). These low-dimensional representations of each subject are then used to generate the weight parameters of an RNN corresponding to the decision-making process of that subject. They applied their model to a synthetic behavior dataset which was generated by simulating a RL model with various parameter values on a two-armed bandit task. The actions taken by the model agents and the received rewards comprised the dataset. Interestingly, it was found that the latent representations have an interpretable relationship with the cognitive parameters that actually determined behavior: the inverse temperature β is mainly related to z_2 and the perseveration parameter κ to z_1 . Another method proposed

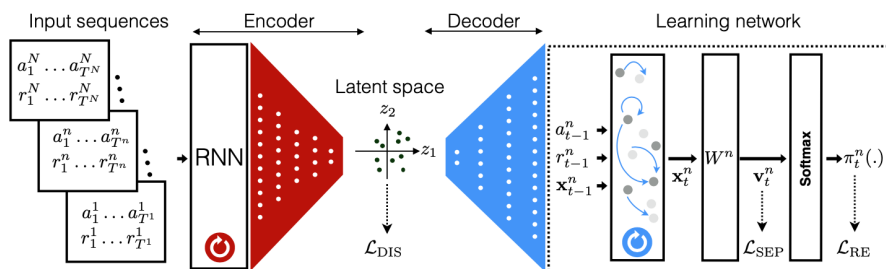


Figure 2.7: The structure of the autoencoder-based framework proposed in [94]. The model comprises an encoder (red) and a decoder (blue). The encoder is the composition of an RNN and a series of fully-connected layers. The encoder maps each whole input sequence (in the rectangles on the left) into a point in the latent space (depicted by (z_1, z_2) -coordinates in the middle) based on the final state of the RNN. The latent representation for each input sequence is then fed into the decoder. The decoder generates the weight parameters of an RNN (called the learning network here) and is shown by the dotted lines on the right side. The learning network is the reconstruction of the closed-loop dynamical process of the subject which generated the corresponding input sequence based on experience. This learning network has similar structure shown in Figure 2.6.

to improve the interpretability is to combine the merits of the data-driven DNN models and the theory-driven cognitive models, that is using DNN model as an exploratory tool to predict human behavior, while using the theory-driven models to explain the DL models [95].

2.7 Conclusion

In this chapter, we firstly defined the two methodologies, i.e. the field experiments and the EMA methodology, adopted in this thesis. The incorporation of behavioral paradigms in these experimental methodologies allows us to explore the manifestations of the computational phenotypes in naturalistic settings. The second section discussed the rationale of using smartphones for using smartphones to implement the two methodologies. The third section introduced the economic assumptions that explain how people make decisions in uncertain and risk conditions and provided evidence in terms of how the neurons implement the theoretical models in the brain. The RL theories and choice models were then delineated in the fourth section to explain the learning mechanisms in paradigms where decision is based on learning from experience. We also reviewed evidence showing how the prediction error in the RL model is expressed in the dopamine neurons. Finally, an

alternative method using RNNs to represent the learning and decision-making strategies used by humans was reviewed.

Chapter 3

Tools for Data Collection and Analysis

This chapter introduces the hardware tools and software applications for data collection in naturalistic settings. The main device we rely on is smartphones, although the web-based and mobile applications also work on tablets or even laptops. Two web-based behavioral paradigms and a mobile application were developed for repeatedly sampling human decision-making in daily routines. The workflow of the hierarchical Bayesian modeling process is described in the last section of this chapter, providing the basic overview of the data analysis methodology adopted in the following case studies.

3.1 The Web-based Decision-making Paradigms

3.1.1 Introduction

Various types of decision-making tasks have been employed to examine the cognitive processes and neural basis underlying decision-making in humans and animals. They originated from different fields including psychology, cognitive neuroscience, and economics. There is also a continual invention of novel paradigms. Some of them are stimulated by the questions generated in the process of computational modelling, while some of them are motivated by the findings from neuroimaging studies in which specific neuronal systems or brain areas are suspected of playing an important role in decision-making. The structure and scheme design of a decision-making paradigm is important for a research question

since the computational models used for describing how information is processed behind the scenes are fundamentally limited by the behavioral data, which is itself in turn limited by the experimental protocol.

We are interested in understanding the between- and within-individual differences of learning and decision-making in contexts with different sources of uncertainties, which are common scenarios in daily life. The structure and scheme of the Iowa Gambling Task (IGT) and the Aversive Learning Task (ALT) are suitable for achieving this purpose. These two decision-making tasks are both designed such that subjects have uncertainty about the consequences of choosing one alternative over the others. They are rich for theoretical modelling of the neural representation, computation, and utilization of different forms of uncertainty that arise during learning and decision-making. The IGT is used to test how people deal with the uncertainty existing in the relationships between actions and outcomes, whereas the ALT is used to examine whether people are able to recognize a higher-order of uncertainty, i.e. the switch of the action-outcome contingencies.

3.1.2 The Web-based Iowa Gambling Task

The IGT is one of the most popular laboratory experimental paradigms for assessing complex, experience-based decision-making across groups. Originally, the IGT was designed to measure the decision-making impairment of patients with lesions in the vmPFC [96]. Since this initial study, this task has been applied to a variety of clinical populations, including patients with neurological disorders [97]–[99] and those with psychiatric disorders [100]–[102]. The IGT is used to measure risk-taking behaviors of patients with chronic pain in the first case study.

Table 3.1: Summary of the payoff scheme of the IGT

	Deck A Bad deck with frequent losses	Deck B Bad deck with infrequent losses	Deck C Good deck with infrequent losses	Deck D Good deck with frequent losses
Reward/trial	100	100	50	50
Number of losses/10 cards	5	1	5	1
Loss/10 cards	-1250	-1250	-250	-250
Net outcome/10 cards	-250	-250	250	250




Figure 3.1: The screenshot of the web-based IGT.

Participants in the IGT are initially given €2000 virtual money and presented with four decks of cards labelled A, B, C, and D. Each card in these decks can generate gains, and sometimes cause losses. Participants have to choose one card from these four decks consecutively until the task shuts off automatically after 100 trials. In each trial, feedback on rewards and losses of their choice and the running tally over all trials so far are given to the participants, but no information is given regarding how many trials they will play and how many trials they have completed during the task. Participants are instructed that they can choose cards from any deck and they can switch decks at any time. They are also told to make as much money as possible minimising losses. Table 3.1 shows the payoffs of the four decks. As can be seen in the table, Decks A and B are two bad decks that generate high immediate, constant rewards, but even higher unpredictable, occasional losses. Thus, the long-term net outcome associated with Deck A and Deck B is negative. In contrast, Decks C and D are two good decks that generate low immediate, constant rewards, but even lower unpredictable, occasional losses. Thus, the long-term net outcome associated with Decks C and D is positive. In addition to the payoff magnitudes, the four decks also differ in the frequency of losses, i.e. Deck A and C are associated with a higher frequency of losses, while Deck B and D are associated with a lower frequency of losses. The key to getting a higher long-term net outcome in this task is to explore all of the decks in the initial stage and then exploit the two good decks. In order to recruit internet participants, we developed a web-based version of IGT (see Figure 3.1 for the screenshot). This implementation is developed using HTML, CSS and Javascript. The server is run on a Heroku instance with the data stored on a managed MongoDB service.

3.1.3 The Web-based ALT

In 2007, Behrens et al. [103] developed a reward gain version of a probability-tracking task in which the volatility of the reward environment was manipulated by switching the action-outcome contingencies, i.e. action-outcome contingencies were stable in one period of the task and volatile in the other. Participants in this task were required to choose between blue and green cards based on both the past success and the reward magnitude associated with each color. According to the RL theory that has been introduced in **Chapter 2**, the learning process of an organism is driven by prediction error, i.e. the difference between expected and actual outcomes. The learning rate that scales the prediction error is a computational phenotype of the participant reflecting the rate at which new information replaces old before it is used to update the participant's belief about the environment. It was argued by Behrens et al. [103] that the learning rate should depend on the current levels of uncertainty in the estimate of the action's value and this uncertainty is determined by the statistics of the reward environment. Thus, the rationale for designing this task is to investigate whether humans can track the statistics of a rewarding environment. Browning et al. [104] modified the task into an aversive version in order to investigate the aversive behavior of anxious individuals. In aversive learning, an aversion is generally created toward a targeted behavior by pairing it with an unpleasant stimulus, such as a painful electric shock or money loss. The two stimuli in the aversive task [104] were associated with the potential receipt of electrical shock. It was thus used to examine if anxious individuals have difficulties in adjusting their learning rates to reflect the stability or volatility of the environment while dealing with negative outcomes. Although the learning rate is a fundamental feature of the behaviour of all organisms and even artificial agents, our understanding to how and why it changes is still less well established. A range of variants of either reward and/or punishment-based learning tasks have been utilized to examine if healthy or clinical populations are capable of maintaining estimates of the uncertainty of the associations they are learning and use these estimations to tune their learning [105]–[107].

We developed a web-based reward-loss version of this task. To maintain the interest and




Figure 3.2: The screenshots of the web-based ALT. In the first trial (panel A), the subject chose the blue leprechaun, which did not steal gold coins from the subject, thus the feedback is "Good choice". In the second trial (panel B), the subject chose the red leprechaun and it stole 36 gold coins and ran away.




Figure 3.3: Outcome probabilities across the course of the ALT. The task comprised two blocks. In the stable block, one leprechaun (for example, the blue one) had a 75% probability of resulting in loss of gold coins; the other shape had a 25% probability of resulting in coin loss. In the volatile block, the probability that the choice of a given leprechaun would result in coin loss switched every 30 trials between 80% and 20%.

enthusiasm of the participants, the task was modified into a more entertaining, leprechaun-themed game. The participants are initially given 10,000 gold coins in the forest. They have to navigate through a clump of bushes to reach home. As they make their way through the bushes, they will come upon a series of junctions. Participants have to repeatedly choose one of two leprechauns, blue or red, to pass through the junctions. At each junction, there will be two leprechauns, blue and red. One of the leprechauns has a higher probability of stealing gold coins from the participants while the other one has a lower probability. Each leprechaun holds a bag in their hands with a number on it, representing how many gold coins the participants will lose if that leprechaun steals from them (see Figure 3.2 for a screenshot of the task. The implementation technology stack is the same as that for the Iowa Gambling Task.). The magnitude values of loss for the blue leprechaun are randomly generated between 1 and 100 (M_{blue}), and the values for the red leprechaun are set to $(100 - M_{blue})$. The participants are instructed that the chance of the stealing leprechaun being blue or red depends only on the recent outcome history. However, because of the loss magnitudes associated with the two leprechauns, participants normally do not select the less likely to steal leprechaun if it is associated with a very high loss. There are two blocks in the task, i.e. the stable block and the volatile block, each for 120 trials. In the stable block, the blue leprechaun has a 75% probability of stealing coins from the participants if it is chosen. The other leprechaun has a 25% probability of stealing. In the volatile block, the probability that the choice of a given leprechaun would result in stealing switches every 20 trials between 80% and 20%. The outcome probabilities of the blue leprechaun across the course of the task are shown in Figure 3.3. The aim of the participants is to choose the better leprechaun that steals less from them as much as possible to reserve more gold coins when they get back home.

3.2 EMA Platform: AthenaCX

3.2.1 Introduction

The implementation of EMA studies is in fact technically intensive. The utilization of AthenaCX, a custom-programmed mobile app, allows us to address and reduce many of the barriers we faced when conducting EMA studies. AthenaCX is not designed only specific for a particular study. It is a platform that enables researchers to build complex EMA investigations. The app has been made available publicly both on the Google Play Store and the App Store. The key features of AthenaCX are described in the following.

3.2.2 E-consent and Survey Design

The Design Studio (DS) (see Figure 3.4 for the home page) is the survey builder where researchers build e-consent, design, configure and deploy digital surveys with different question types in AthenaCX. In order to implement their experiments, researchers can start by creating e-consent in the DS. AthenaCX enables them to easily create and distribute consent surveys that stand up to Ethics Committee Review Boards scrutiny by implementing a thorough and robust 'staged consent' flow. It works well for studies being carried out in both an online and offline capacity. After the e-consent, researchers can then start creating surveys in the DS. They can have as many as surveys as they like within a particular study by adding survey cards. Multiple question types are available in the DS (Figure 3.5), including:

- Verbal Response Scale (VRS). Participants answer by selecting one of the predefined options you create and list vertically.
- Visual Analogue Scale (VAS). Provide participants with an on-screen scale with which they can provide their answers through the positioning of a handle, which is commonly referred to as a 'slider'.
- Numeric Response Scale (NRS). Provide participants with a set number of predefined boxes which are representations of a numerical answer. A common use is to

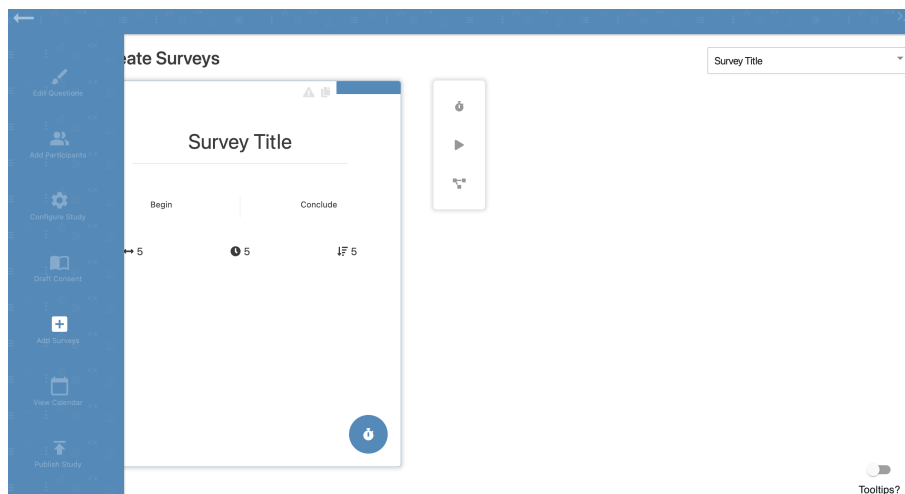


Figure 3.4: The home page screenshot of the DS.

provide an answer using the numbers 0 - 10.

- Text Box. Participants give their answers through a free text field.
- DateTime Widget. Allow participants to answer in the format of a date, time or date and time.
- Instructional Widget. Provide participants with information through the message created here. This is useful for generation of ecological momentary interventions

Powerful branching logic has been embedded in AthenaCX, which means researchers can change what questions or surveys a respondent will see based on how they answer a previous question. Branching is compatible with question types with defined answers, such as VRS, VAS, and NRS. This feature makes it possible to show the respondent just the fields they should see and make the sequence of the questions as short as possible, ultimately improving the completion rate of the measurements.

3.2.3 Sampling Schemes

The sampling and assessment scheme is the cornerstone of an EMA procedure. AthenaCX can be customized for various types of designs by using different types of surveys and altering when notifications are scheduled within the configuration page in the DS. Three types of surveys (Figure 3.6), i.e. time-based, self-initiated and remote trigger, and a set of

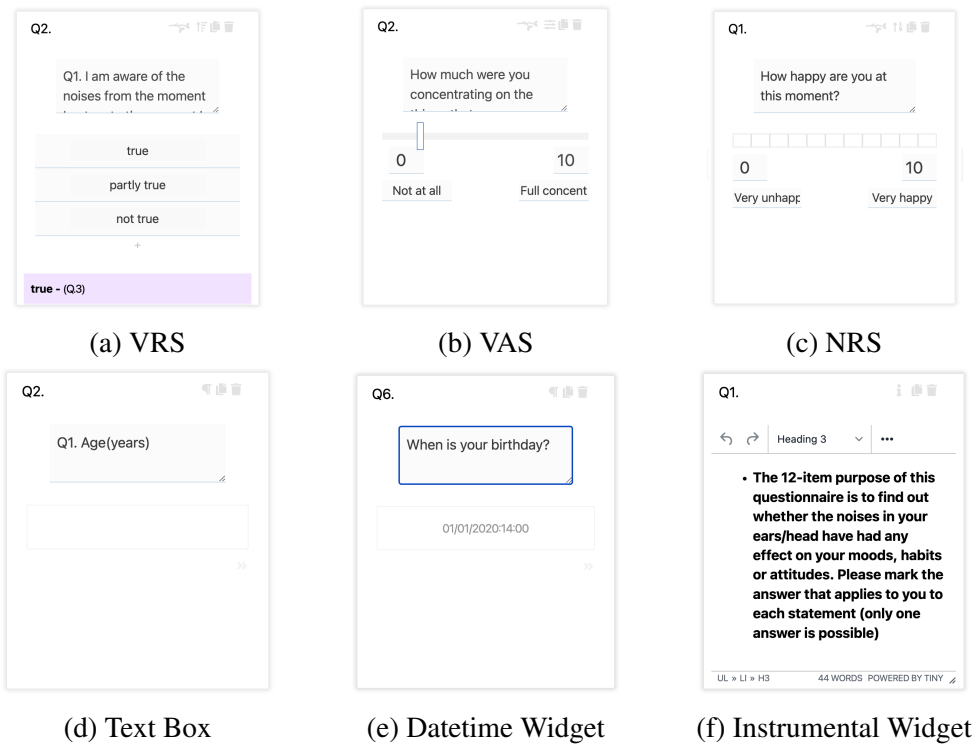


Figure 3.5: Question types provided in AthenaCX.

configuration parameters, e.g. start and end time, the time between the surveys, and active time of the survey etc., are provided in AthenaCX to achieve various sampling schemes.

The EMA sampling and assessment schemes can be roughly divided into event-based sampling and time-based sampling schemes [108]. Event-based schemes aim to sample discrete events or episodes in subjects' lives and schedule the data collection around these events, whereas time-based schemes focus on characterizing subjects' entire experience more broadly and inclusively. Some clinical and research interest is in particular events or episodes, e.g., instances of drinking, smoking, or panic attacks. These cases lend themselves to event-based monitoring, in which data collection is triggered by the occurrence of a predefined event of interest to the investigator. Typically, subjects themselves determine when the event has occurred and initiate an assessment. It is possible to automatically detect some events through incorporating wearable data, e.g. physical activities, sleeping patterns, and mobility patterns collected with Fitbit and other wearables, in AthenaCX. The remote trigger survey enables researchers to trigger and schedule specific surveys based directly on their dynamic health and activity data. If the event can not be detected by the devices, researchers can use the self-initiated survey that is always available to the subjects

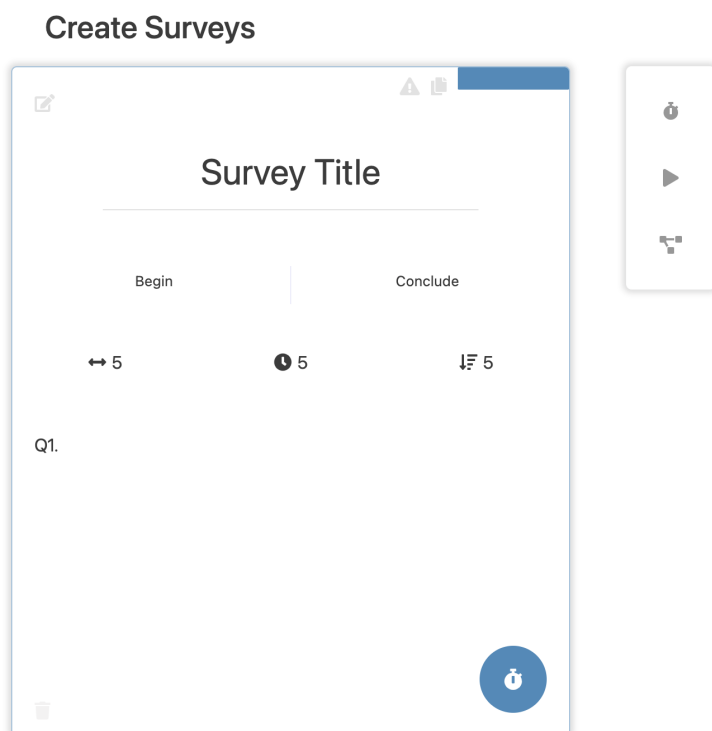


Figure 3.6: Survey types provided in AthenaCX. Inside the panel on the right side of the survey card, the alarm clock is the time-based survey. Researchers can use this to create schedules based on specified timing patterns, e.g. a daily or weekly survey. The triangle is the Self-Initiated survey. This survey type will always be available to the participant during the time period that it has been configured to be available. The third sharing icon is a remote trigger survey. This survey type can be activated or scheduled directly from the dashboard at any point once the participant has enrolled.

during the time period. As a result, subjects can respond to questions in AthenaCX on their own time whenever the event occurs.

Some clinical phenomena, e.g. actigraphy, heart rate, and skin conductance, vary continuously and can be monitored continuously using wearable devices, while many other phenomena, e.g. mood and pain, vary continuously but are not easily conceptualized in an episodic framework either. EMA protocols rely on time-based sampling for these cases. There are many variants of time-based sampling schemes, which vary in schedule, frequency, and timing. A time-based survey in AthenaCX can be used to create this type of schemes based on any specified timing patterns either at fixed intervals or random times. For interval-contingent designs (i.e., participants complete questionnaires at predetermined times), notifications can be set to repeat every day of the experience sampling phase at specified times. For signal-contingent designs (i.e., participants complete questionnaires whenever they are signalled), researchers can set signals to occur at set intervals (e.g., every 2 hours) or randomly within a specified range (e.g., every 1.5 to 3 hours after each signal) in a set period of time (i.e., daily data collection period).

3.2.4 Push Notifications

A push notification is a mobile alert visible on the locked screen of a smartphone or tablet and it is the solution used in AthenaCX to remind the subjects to complete EMA surveys. When it is time to take a survey, a notification pops up and chimes to signal to participants that they should take a survey. Participants tap the notification to launch the app, or they can tap on the application's icon directly to launch it and answer the questions. AthenaCX uses the local notification system available on Android and iOS devices to carry out the signalling function essential for EMA studies. The local notification system is "local" because it does not require a cell signal or an internet connection. Notifications are scheduled immediately after participants have installed and set up the app on their phones. The scheduled notifications are recorded and sent to the server so that researchers can check response latencies and compliance. The local notification system also allows the app to have a snooze function, which enables participants to delay answering the questionnaire

for a specified period of time before they are sent another reminder to complete the survey.

3.2.5 Offline Data Collection

One of the biggest strengths of mobile apps is the possibility to work offline; in other words, without an internet connection. This technology opens the door to EMA studies in places with slow or no internet connection or areas where the cost of mobile internet is too high. AthenaCX supports offline data entry, which means it stores the data in the local memory of the device and synchronizes it to the server whenever an internet connection is available. Each response is time-stamped to the millisecond, so researchers can determine whether responses were actually collected in the same moments, or whether the participant was interrupted in the middle of completing the questionnaire and resumed that questionnaire at a later point in time. When the researchers are trying to identify 'careless EMA responses', the time taken to complete an item can be very useful.

3.2.6 Link with External Web-based Tasks

AthenaCX is valuable for decision-making studies due to its support for incorporating decision-making paradigms that are designed to provide objective measures of behavior, along with self-report EMA surveys. As explained before, the adoption of decision-making paradigms enables us to conduct computational modelling analysis, which offers new insights to the underlying mechanisms. Typically, decision-making paradigms are applied in cross-sectional studies where retrospective symptom inventories are used to predict task performance. However, it is difficult to falsify the impact of the fluctuation of the emotional or physical status of the subjects with cross-sectional methods. Although the most economical and flexible way to collect experience data and task data is with a custom programmed app, designing an EMA app embedded with decision-making tasks from scratch requires advanced programming skills and a large amount of time and energy. AthenaCX is designed to eliminate the barriers of integrating within-subjects task data in longitudinal studies. The researchers only need to implement the behavioral paradigms they wish to use in their experiments themselves. Sometimes, they are able to find the tasks

on open source platforms, e.g. Pavlovia, where hundreds of tasks have been implemented and shared for free research uses. It is simple to configure the URL to their task in the DS with an instructional widget question. The URL will be sent to AthenaCX and AthenaCX will append the participant ID to the URL. The researcher only needs to make sure their web-based application is able to read the URL and extract the participant ID in their source code when the participants are directed to their fault browser. It means it is not necessary for the participants to type in their IDs again even if they are directed outside the AthenaCX, which could be a great burden for them and might impact the compliance rate.

3.2.7 Participants Onboarding

Onboarding participants is a crucial element in EMA studies. Participants can download the AthenaCX mobile app to their phone using rapid onboarding via an email link, QR Code, or app store search. After getting into the app, participants can search within the app among different public studies that have been created by AthenaCX researchers and easily join them. For closed studies that require pre-screening to select participants, researchers can make their studies as closed studies so that they are not available to everyone. They can create customized participant IDs within the DS and then share the IDs with those who are qualified to participate. Participants are required to type in the IDs in AthenaCX to join the closed studies.

3.2.8 Data Security

When doing daily diary studies, sensitive personal data, and possibly even medical data, will be collected. Hence, data security is of significant importance and is the first thing we need to consider. The databases of AthenaCX physically reside on the Amazon Web Services platform in its Western Europe (Ireland) data center and are fully managed by MongoDB Atlas. Data is transferred using the Heroku platform. Heroku provides fully secure services which operate using HTTPS, Secure Sockets Layer (SSL). Heroku uses SSL for all authentication (logins), and administration of the site. The user's browser establishes the authenticity by requesting an SSL certificate that verifies the identity of

AthenaCX. Once that SSL certificate is recognized, an SSL connection is established for security, encrypting data transmitted between browser and web server. Additionally, only the researchers have the access to the data.

3.3 Workflow of Theory-Driven Computational Modeling

3.3.1 Introduction

The rapid advances in neuroscience have established that the brain, the most complex organ, is essentially a massively complex nonlinear computational device that governs the expression of an organism's behavior. Thus, computational depictions of cognitive function have been widely advocated in recent literature. Different from conventional statistical modelling tools that generally ignore theoretically-relevant data-generating mechanisms under behavior, the key feature of applying computational modelling in behavioral and brain science is the definition of a formal mechanistic cognitive process, i.e. using mathematical principles and formalism to generate consistent, rigorous, and testable hypotheses about the underlying mechanisms of the brain and make better sense of the behavioral data [109], [110]. As a result, computational models can link the trial-by-trial experimentally observable behaviors, e.g. stimuli, outcomes, and experiences, to behaviors in the immediate future, while traditional statistical models only provide superficial summary descriptions of observed behaviors, e.g. actions or response time averaged over time or trials. The collection of free parameters derived from computational models fit to behavioral data forms the concept of computational phenotype that has interpretations in terms of cognitive processes. This shifts our focus away from using descriptive trait and symptom measures toward estimating psychologically interpretable parameters to represent psychological characteristics [111]. Furthermore, it was found that the application of computational modelling could produce higher test-retest reliability estimates, alleviating the pressure of the reliability paradox in the study of individual differences and thus advancing theory development [112], [113].

In a standard decision-making paradigm, such as the IGT and the ALT, a participant is

presented with multiple stimuli to choose repeatedly and receives rewards or punishments according to their choice on each trial. While designing the task, a theoretical framework is built for the task with a set of assumptions and mathematical equations. Multiple computational models are then formulated to instantiate the assumptions about how the behavior is generated. The behavioral data is not necessary limited to be choice data, instead, it could be any observable data that can be quantified statistically in terms of the probability the computational models assign to the data, such as reaction times, eye movements, and even neural data. We will be mainly using choice data in this thesis. In spite of the differences in terms of the structures and application situations, the general workflow of examining the explanatory power of such models is the same. Typically, a set of free parameters that allow the model to produce a range of behaviors are specified along with the structural theory assumptions. Once the experimental paradigms and computational models are designed, an important and necessary step, is to conduct model simulation and generate synthetic behavioral data with various parameter values. This step is helpful to refine the experimental and model design before data collection. After collecting real data from the participants, they are fitted to the predefined models and the free parameters of the models are solved through a process known as parameter estimation. Different computational models constitute different hypotheses about the cognitive process that gives rise to the data. In order to test which hypothesis fits the data best, the next critical step is to evaluate and validate the model from different aspects.

3.3.2 Model Simulation

Model simulation involves running the model with known model and known parameter settings to generate synthetic behavioral data. Simulation is important and useful across all phases of computational modelling, no matter before or after collecting real data from the participants [114]. In this section, we mainly focus on the use cases of simulation that aim to improve experimental and model design before data collection. Simulation for further checking the model after parameter estimation of the real data is referred to as model validation and is introduced in the following section. Before the real data arrive,

simulation allows us to understand how behavior changes with different models across the range of parameter values and refine the prior distribution of the parameters. To simulate behavioral data, we need to first determine the task parameters, e.g. the number of trials, the reward or loss magnitudes and the probability of reward or loss for each option across trials, which should match the actual parameters used in the experiment. We can then randomly sample a set of free parameters $\tilde{\theta}$ from the prior distributions $\pi(\theta)$ over each parameter. With the parameter set $\tilde{\theta}$, a sequence of behavioral choices \tilde{y} can be generated from the model: $\tilde{y} \sim \pi(y|\theta)$. This process can be repeated as many times as we wish with different parameter settings, yielding a matrix of simulated choices.

In order to assess whether the simulated choices capture key aspects of behavioral patterns of interest and whether the prior model predictions are consistent with domain expertise, it is useful to define model-independent summary statistics to visualize the simulated data. The model-independent measures provide key markers of what we want the model to account for in the data. They should thus be carefully chosen and designed based on expectations we have about the true data generating process and about the kinds of structures and effects we expect the data may exhibit. For instance, the proportion of choosing good decks is a widely used model-independent measure for the IGT. The model-independent measures for the decision-making tasks were suggested to be summarized at three levels [115]. At the trial level, the model-independent measures are used to examine the trial-by-trial dynamic of the choice behavior averaging across participants. At the participant level, model-independent measures can be used to assess the individual variation averaging across trials. At the overall level, model-independent measures provide the average performance of the model across both trials and participants. To draw the inference, correlation analysis can be conducted to determine how the measure-independent statistics change with different parameter settings.

3.3.3 Parameter Estimation

Parameter estimation is a process that tries to minimize the error between the model's predictions achieved with a particular combination of parameters and the empirical data.

Typically, a behavioral dataset includes a number of subjects, each with a series of choices on the decision-making task. There are a number of different ways for estimating the parameters. Traditionally, parameters are estimated with the most widely used frequentist approach, Maximum Likelihood Estimation (MLE) [116] either at the individual level or group level. The MLE is a point estimate based on the likelihood function alone. It assumes there is a set of unknown but fixed parameters θ and estimates θ with some confidence. Group-level MLE simply aggregates likelihood across both trials and subjects and estimates a single set of parameters for the whole group of individuals. This method treats the data from all subjects as though they were just more data from one individual. The estimated parameters do not vary across subjects, thus individual differences can not be derived, which is unlikely to be the case in reality. Individual-level MLE estimates a set of maximum likelihood parameters for each individual. Although individual differences are considered, estimates of this method are often noisy and unreliable, easily leading to extremely parameter values, especially for psychological studies with few observations for each participant.

It is thus more advocated to utilize hierarchical models [117]–[120] in which individual differences and commonalities are both considered. Hierarchical models assume individual parameters are drawn from a common distribution defined by group-level parameters and the group and individual level parameters are estimated simultaneously. The accuracy and precision of parameter estimation can be improved through the shrinkage effect [121], [122] of hierarchical models, i.e. data is pooled across individuals to explicitly estimate the mean and variance of the population and the group estimates then in turn inform the parameter estimation of each individual in the group, thus avoiding unlikely extreme values. However, hierarchical models are conceptually and computationally inconvenient to implement in a frequentist context but are conceptually simple and computationally tractable in Bayesian frameworks. Another benefit of Bayesian analysis is that it could provide full posterior distribution to quantify uncertainty about the free parameters, whereas MLE only generates fixed point estimates.

As a result, the combination of the concept of hierarchical modelling and the Bayesian

framework yields probably the most competitive parameter estimation approach, i.e. Hierarchical Bayesian Modeling (HBM), in the literature [122]–[124]. Recent years have witnessed an emerging rise in the use of Bayesian analysis in cognitive science due to its increasing recognition advantages in terms of robustness and reliability. The Bayesian posteriors can be approximated with various methods. The most prominent class of methods to approximate Bayesian posteriors are Markov Chain Monte Carlo (MCMC) methods, which are very technical and complicated to implement. However, the increasing availability of specialized programming languages and packages make these previously less accessible algorithms easy to use. Several probabilistic programming languages, such as Stan [125], WinBugs [126], and JAGS [127] have been developed, in which MCMC algorithms are embedded for automatically generating MCMC samples. Many of these MCMC generators use the corresponding language to define the model and rely on a different programming language (like R, Python, Julia, etc) to communicate with the program that does the sampling. Additionally, a series of tools and packages including the hBayesDM package, the VBA toolbox, the HGF toolbox in the TAPAS software collection, the *brm* and the *rstanarm package*, etc, requiring minimal knowledge of programming to fit complex hierarchical Bayesian models are also available.

3.3.4 Model Comparison

In model comparison, our goal is to determine which model, out of a set of plausible models, is most likely to explain the actual data. It is known that adding more parameters nearly always improves the fit of a model to the sample. Here, 'fit' means a measure of how well a model can retrodict the data used to fit the model. Nevertheless, the problem with more complex models is they often predict new data worse, a problem termed overfitting. Specifically, a more complex model will be very sensitive to the exact sample used to fit it, leading to potentially large mistakes when future data is not exactly like the past data. In contrast, simple models with too few parameters, tend instead to underfit, systematically over-predicting or under-predicting the data, regardless of how well future data resemble past data. As a result, we cannot favour either simple models or complex models only based

on their within sample predictive accuracy. The best practice to cope with the trade-off between goodness-of-fit and complexity would be testing the model on out-of-sample data. However, this is difficult to achieve in practice, especially given the fact that the sample size in cognitive studies is typically small. Another challenge is that in order to measure the predictive performance of the models, we generally have to know the target model, which is usually not accessible. We would not be doing statistical inference if the target model is already known. A series of strategies have been developed to tackle these problems. The most popular two strategies were information criteria and cross-validation, both of which provided a relative measure estimating the predictive accuracy of different candidates without knowing the target model. It is worth noting that since the two strategies are relative measures, we can only use them to estimate which candidate model gets closer and by how much to the target model, but we would not know how far each model is from the target model.

Various information criteria have been adopted in literature including Akaike Information Criterion (AIC) [128], Deviance Information Criterion (DIC) [129], Bayesian Information Criterion [130], and Widely Applicable Information Criterion (WAIC) [131]. Essentially, all of them apply a correction to correct for the selection bias caused by fitting and validating the model on the same sample of data, aiming to come out with an estimated measure of out-of-sample predictive accuracy. Although each of them designed the measure with different methods, the basic idea of calculating these criteria is the same, i.e. evaluating the fit of the model on the data that was used to train the model and then subtracting away a penalty term which accounts for the selection bias. Common to all of these criteria, they evaluate the fit using log-likelihood and the penalty basically measures the degree of uncertainty in the parameters. The measure of fit which AIC uses is the log-likelihood evaluated at the maximum likelihood estimate of the parameters and the penalty that AIC applies to the with-in sample fit is simply the number of parameters in the model. It is of mainly historical interest now because the theoretical basis of AIC is not applicable in more complex context. DIC uses a similar measure of fit at a point estimate of the parameter value. Instead of the log of likelihood evaluated at the maximum likelihood

estimate, DIC uses the log-likelihood evaluated at the maximum posterior density of the parameters. A big change with DIC is the application of a more general and useful measure of penalty, which is 2 times the variance across all of posterior draws. WAIC is a criterion that is more general than both AIC and DIC. Instead of using the point estimate, WAIC calculates the average log-likelihood over all of the posterior samples. The penalty part of WAIC is also evaluated pointwise, computing the variance in log-likelihood for each observation in posterior distribution.

Another popular strategy is to actually use out-of-sample data to measure the fits of the models. This is known as cross-validation, dividing the sample in a number of chunks (or folds) and predicting each fold after training on all others. The extreme case of cross-validation is to take each unique observation as a fold, leading to the maximum number of folds, which is called leave-one-out cross-validation (LOOCV). The key trouble with LOOCV is that, if we have 1000 observations, that means computing 1000 posterior distributions, which can be time-consuming. There are clever ways to approximate the LOOCV without actually running the model over and over again. One way is to take the pointwise log-likelihood as an estimate of the out-of-sample fit. Additionally, the adoption of importance sampling, e.g. Pareto-smoothed Importance Sampling (PSIS), provides an even faster and more reliable approximation of LOOCV predictive accuracy, i.e. PSIS-LOO. PSIS-LOO and WAIC will be adopted in this thesis for model comparison because these two indices use the pointwise log likelihood of the full Bayesian posterior distribution instead of point estimates. This is compatible with the HBM methodology we are utilizing for parameter estimation. Similar to the WAIC, approximation of PSIS-LOO is also on the information criterion scale; thus, lower values of these indices indicate better out-of-sample prediction accuracy of the candidate model. Both of them can be computed using an R package called *loo* [132].

3.3.5 Posterior Predictive Check

Sometimes, it is not enough to select the model with the lowest criterion value of either PSIS-LOO or WAIC because both of them are only relative comparison criteria that merely

consider the merits of each model's performance relative to the rest. It means they imply no absolute criterion for model selection. Thus, there is no guarantee that the winning model survived from these indices indeed explains or predicts the behavioral effect of interest. It was suggested to perform other model comparisons in absolute scale in-between deciding on the winning model.

One of the most widely used model validation approaches is posterior predictive check (PPC) [122], [133], [134]. To generate the data used for PPC, we can simulate from the posterior predictive distribution. The posterior predictive distribution is the distribution of the outcome variable implied by a model after using the observed data y (the shape of which is $N \times T$, where N is the number of subjects and T is the number of trials) to update our beliefs about the unknown parameters θ in the model. For each draw of the parameters θ from the posterior distribution $p(\theta|y)$, we generate an entire matrix of outcomes. The result y_{rep} is an $S \times N \times T$ array of simulations, where S is the size of the posterior sample. That is, each $N \times T$ matrix is an individual "replicated" dataset of the observed data y :

$$p(y_{rep}|y) = \int p(y_{rep}|\theta)p(\theta|y)d\theta \quad (3.1)$$

where

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)} \propto p(y|\theta)p(\theta) \quad (3.2)$$

The predictive density, $p(y_{rep}|y)$, depicts the degree to which model-reproduced data (y_{rep}) corresponds to the actual data (y). It is worth noting that the generation of this replicated dataset y_{rep} is conditioned on the observed response data. In order to quantify how much the replicated predictions of a model deviate from the actual data, we can use the model-independent measures introduced in previous section to summarize both the observed data and the replicated data at different levels. At either trial level or individual level, we can examine the correlations of the model-independent measures between the replicated data and the observed data. At the overall level, a Bayesian p-value [122] could be computed to assess how much area under the posterior curve is below the actual data.

3.3.6 Parameter Recovery

Apart from assessing the predictive performance of the models, another assessment is to evaluate how well a model can recover known parameter values. This is the process of parameter recovery where the relationship between true (simulated) and recovered parameters is examined [84]. In order to do this, we need to choose a set of ground truth parameters first. Generally, the ground truth parameter set is generated by taking the means of individual-level posterior distributions of the model to fit the real data to ensure that the true parameter values were reasonably distributed and representative of human decision makers for the corresponding model. Using the selected true parameters, we are able to simulate choice datasets for a given model and task structure. The task structure is exactly the same as that used in real experiments. The model is then fit to the simulated datasets and the recovered parameter estimates are compared to the true parameters. According to [135], we can either use the means of the posterior distribution or the entire posterior distribution of each individual-level parameter to infer how well the model could recover parameters in an absolute and relative sense.

3.4 Conclusion

In this chapter, we introduced the origins and implementations of the web-based behavioral paradigms adopted for sensing human decision-making in naturalistic contexts. We then presented the mobile application, AthenaCX, which was used to configure and design complex EMA sampling schemes in our studies. Lastly, the workflow of the computational modeling analysis for the trial-by-trial data of the decision-making paradigms is presented. The principal tools and workflows set up here in this chapter lay a solid foundation for the case studies in the following chapters.

Chapter 4

Differences in Reward and Persistency Characterize Behavior in Chronic Pain for the IGT: A Lab-in-the-field Study

In this chapter, we conducted a lab-in-the-field experiment to capture how chronic pain impact decision-making on the IGT. A cross-sectional design with independent groups was employed. A convenience sample of 45 participants were recruited through social media, including 20 participants who self-reported living with chronic pain, and 25 people with no pain or who were living with pain for less than 6 months acting as controls. All participants completed both pain experience assessments and the IGT in their web browser at a time and location of their choice without supervision. Standard behavioral analysis revealed no differences in learning strategies between the two groups, computational modelling analysis suggested that compared to healthy controls, individuals with chronic pain were more driven by rewards and less consistent while making choices between the four decks of IGT. This lab-in-the-field case study demonstrated that 1) It is feasible to administer the behavioral paradigms in the naturalistic environments and significant effect of the variables of interest could be detected even though there is a high chance that additional noise was included in such settings; 2) compared to standard statistical summaries of behavior performance, computational approaches offered superior ability to

resolve, understand and explain the differences in decision-making behavior in the context of chronic pain outside the laboratory.

4.1 Introduction

Chronic pain is defined as pain persisting or reoccurring for more than three to six months [136] and has been recognized as one of the most significant health issues of the 21st century [137], [138]. Approximately 100 million adults in the US experience chronic pain, resulting in an annual cost of \$560-\$635 billion in medical treatment and lost productivity [139]. Globally, it is estimated that around 20% of the population lives with chronic pain [140], [141]. As a result, there is significant ongoing research into understanding chronic pain and supporting people that live with the condition. One key area of research is the impact of chronic pain on cognitive or neuropsychological abilities. It has been reported that at least 20% of clinical patients with chronic pain, including those without a history of mental disorders, complain of cognitive impairments that cause significant difficulties to their social life and daily functioning [142]. In other studies, researchers have found that cognitive deficits occur across a range of pain conditions, including fibromyalgia [143], migraine [144], chronic back pain [145], and chronic neuropathic pain [146].

Although pain is an attention-demanding sensory process, cognitive alterations cannot be simply attributed to the extra attentional demand from ongoing pain [145]. fMRI has revealed decreased grey matter density in the vmPFC area [147]–[149] and less brain activation in cortical structures during response inhibition among patients with chronic back pain [150]. These findings are important as the vmPFC also plays a critical role in other cognitive functions, such as decision-making [151], executive control [152], learning [153], and memory [154]. The latest research findings have confirmed that reduced glutamate in the mPFC significantly impairs emotional and cognitive processing in people with chronic pain [155]. This suggests that chronic pain may have a negative impact on the vmPFC and related neural structures and could be considered a cognitive state that may be competing with other cognitive abilities, especially those involving the vmPFC, such as decision-making, which is one of the cognitive domains in which individual with chronic

pain are commonly impaired.

The IGT developed by Bechara et al [96] has been used for classifying various clinical populations from healthy controls, such as patients with lesions to the vmPFC [156], [157], obsessive-compulsive disorder [158], and chronic cannabis use [118], [159]. It has also been applied to investigate abnormalities in decision-making among people living with chronic pain conditions, yielding significant findings. After extracting behavioral responses to the IGT, Apkarian et al. [147] found that patients with chronic pain chose cards from bad decks more frequently. The gambling performance on the IGT was found negatively correlated with the reported intensity of chronic pain. Subjects in [160] were required to complete both the original IGT, where the reward is immediate and punishment is delayed and an IGT variant, where punishment is immediate and reward is delayed. The authors summarized their behavioral choices and found that women with Fibromyalgia had significantly lower scores on the third block, which was referred to as the hunch period of the task, on the original IGT, while intact performance on the IGT variant, suggested these patients were hypersensitive to rewards. Similar results were obtained in [145] where people with chronic back pain won significantly less money relative to healthy controls and their IGT scores did not change significantly throughout the whole task.

The alterations in decision-making on the IGT have been linked to changes in brain function and structure and they may serve as potential biomarkers, aiding the development of more accurate diagnostic tools and treatments for patients with chronic pain. However, all the relevant studies were conducted in laboratory settings where the subjects were under tight experimental control. No study has been conducted to examine whether these results can be replicated when subjected to real-world settings, which is an important question need to be addressed for clinical translation. We aim to measure subjects' decision-making performance on the IGT in naturalistic settings. Additionally, pushing these objective measures of decision-making outside of the lab is a crucial step to translate lab-based research findings into clinical settings and makes it possible to use it as a daily self-checking tool for the patients. Thus, the lab-in-the-field methodology is adopted to investigate the differences in characterizations of decision-making between individuals with chronic pain

and healthy controls in their everyday living environment.

Another limitation in the literature of using the IGT to investigate chronic pain was that the computational modelling analysis was rarely applied. We only found one reference in which a simple heuristic model was designed to differentiate the behavioral performance of individuals with chronic pain and healthy controls on the IGT [161], thus the changes in latent cognitive parameters that drive the impaired performances of people living with chronic pain remain unexplored until now. In this study, we applied computational modelling analysis as a complementary approach to explicitly decompose behavioral performance on the IGT down into cognitive parameters. It has been documented that estimated parameters from such models are able to reveal group differences in cognitive processes despite an absence of group differences in conventional IGT measurement [162]. In summary, this case study is among the first research studies investigating the alterations of the computational phenotypes that characterize the decision-making of the chronic pain population through an online lab-in-the-field experiment.

4.2 Methods

4.2.1 Recruitment

The subjects in this study were recruited through social media and local pain advocacy groups. A total of 64 people, including 28 symptomatic participants who have lived with chronic pain for months and 36 healthy controls who have never lived with chronic or have suffered pain for less than 6 months, were interested in participating in the experiments. They completed the questionnaires and the IGT through computer or mobile phone. 8 symptomatic participants and 11 healthy controls were excluded from the study because they failed to finish the whole IGT. Thus, a convenience sample of 45 participants (70.3% of the total) was recruited finally, consisting of $N = 20$ symptomatic participants (14 females, age = 40 ± 12), and $N = 25$ healthy controls (12 females, age = 38 ± 12). The groups did not significantly differ in age ($P = 0.43$) and the proportion of females ($P = 0.07$). The study was approved by the local ethics committee of the School of Computing, Dublin City

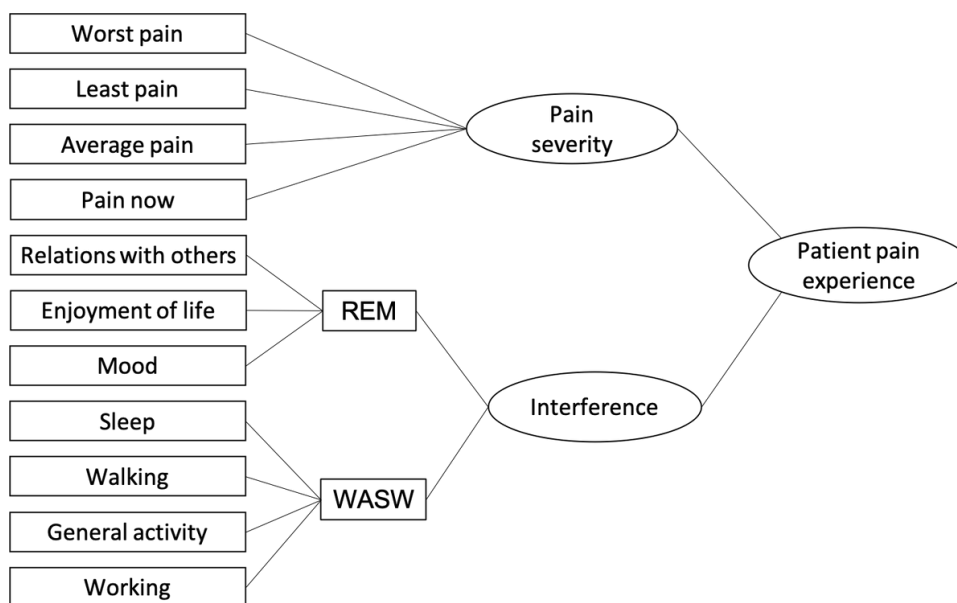


Figure 4.1: A graphical representation of the conceptual framework of the measurement.

University, since it is notification only study. Participants were provided a plain language statement and informed consent when they were involved.

4.2.2 Assessment of Pain Experience

The Brief Pain Inventory (BPI) - Short Form is a validated 9-item self-administered questionnaire used to evaluate the severity of the patient's pain and the impact of the pain on the patient's daily functioning. The patient is asked to rate their worst, least, average, and current pain intensity, list current treatments and their perceived effectiveness and rate the degree that pain interferes with general activity, mood, walking ability, normal work, relations with other persons, sleep, and enjoyment of life on a 10-point scale. The pain interferences are divided into the affective subdimension (i.e. relations with others, enjoyment of life, and mood (REM)) and activity subdimension (i.e. walking, general activity, sleep and work (WASW)). The BPI-Short Form has been used with a variety of populations and across these studies has been shown to be a valid and reliable measure, with adequate internal reliability (e.g., $\alpha = .86 - .96$) [163]. A graphical representation of the conceptual framework of the measurement is shown in Figure 4.1.

4.2.3 The Web-based IGT

Participants were given €2000 virtual money at the beginning and were asked to choose between the four decks of cards for 100 trials in the web-based IGT. The payoff schemes and more details of the task have been introduced in **Chapter 3** and we will not elaborate further on the implementation here.

4.2.4 Standard Behavioral Data Analysis

To compare the behavioral differences between the two groups on the IGT, three parameters were measured. Firstly, the total amount of gain at the end of the task was calculated for each participant to measure their overall performance on this task. An unpaired t-test was used to determine if the difference between the two groups was significantly different on this measure. Secondly, in order to obtain a visual exploration of the group-level deck preferences across trials, we calculated the proportions of good deck selections (Deck C and Deck D) and the learning IGT scores (i.e. the difference between the number of good deck selections and the number of bad deck selections) across the task. Specifically, five new variables were created through the division of the 100 trials into 5 blocks of 20 trials each without overlap. The proportion of good deck selections and the learning score in each block were calculated. In this way, five proportions and five learning scores, one for each block, were obtained for each participant. The comparison between the five learning scores was regarded as an index of learning. A learning score increasing from the first block to the last block suggests that a subject is developing a preference for good decks and an effective selection strategy. Given the repeated learning scores of the two groups over five blocks of trials, a block-by-group repeated measures ANOVA analysis was performed (within-subjects factor: block 1-5, between-subject factor: group healthy vs chronic pain) to reveal whether the two groups differ in learning curves.

4.2.5 Computational Modelling Analysis

Multiple cognitive models have been proposed to identify the psychological processes that drive participants' behavioral performances on the IGT. The most promising models of the IGT according to the literature [164], [165] are the Expectancy-Valence Learning (EVL) model (which is also the first proposed cognitive model for the IGT) [166], the Prospect Valence Learning model with Delta (PVL-Delta) [164], the Prospect Valence Learning model with decay (PVL-decay) [118] and the Values-Plus-Perseverance (VPP) model [118], and the latest proposed Outcome-Representation Learning (ORL) model [135]. We fitted the four models for the two groups separately using the hBayesDM package for R [167], in which HBM for the IGT was implemented.

The PVL models. Both the PVL-delta and the PVL-decay model consist of a utility function, a learning rule, and an action selection rule. They are identical except for using different learning rules. The utility function determines the weight given to gains relative to losses. Both of the two PVL variants applied the utility function from the Prospect Theory [64] that featured diminishing sensitivity to increases in magnitude, and different sensitivity to losses vs. gains. The utility $u(t)$ of each net outcome $x(t)$ (i.e. the difference between the amount of rewards and losses) on trial t is calculated as follows:

$$u(t) = \begin{cases} x(t)^\alpha, & \text{if } x(t) \geq 0 \\ -\lambda |x(t)|^\alpha, & \text{if } x(t) < 0 \end{cases} \quad (4.1)$$

in which $u(t)$ is the subjective utility of the experienced net outcome $x(t)$; α is the outcome sensitivity parameter ($0 < \alpha < 1$) that controls the shape of the utility function and λ is the loss aversion parameter ($0 < \lambda < 10$) that governs the sensitivity to losses relative to gains. Individuals with a higher value of α suggest they have greater sensitivity to feedback outcomes. Individuals with a value of λ less than 1 indicate that they are more sensitive to gains than to losses and those with a value of λ greater than 1 indicate that they are more sensitive to losses than to gains.

The learning rule in the PVL models is used to update the expectancies of the decks $E(t)$ based on the subjective utility value. In the delta rule, a simplified version of the RW

rule is applied, in which only the expectancy of the chosen selection is updated, while the expectancies for other decks remain unchanged:

$$E_i(t+1) = E_i(t) + A \cdot (u(t) - E_i(t)) \quad (4.2)$$

in which A is the learning rate parameter ($0 \leq A \leq 1$) that determines how much weight the decision maker gives to the recent outcomes when updating expectancies. In the decay rule, however, the expectancies of all decks are discounted on each trial except for the chosen deck, which is updated by the current outcome utility:

$$E_i(t+1) = A \cdot E_i(t) + \delta_i \cdot u(t) \quad (4.3)$$

here, A is the decay parameter ($0 \leq A \leq 1$) that determines how much the past expectancy is discounted and $\delta_j(t)$ is a dummy variable which is 1 when deck i is chosen and 0 otherwise.

The action-selection rule generates the choice possibilities $P(D(t+1)) = i$ for each deck for the next trial using a softmax function:

$$P(D(t+1) = i) = \frac{e^{\theta \cdot E_i(t+1)}}{\sum_{k=1}^4 e^{\theta \cdot E_k(t+1)}} \quad (4.4)$$

where $D(t)$ is the chosen deck on trial t , and θ is assumed to be trial-independent and set to $3^c - 1$, and c ($0 \leq c \leq 5$) is a response consistency parameter. A higher value of c represents that the decision maker has a higher tendency to select choices with higher expected values, which means they are responding more deterministically.

The VPP model. The VPP model adds a perseveration term $P_i(t)$ for the chosen deck i on trial t based on the PVL-delta model:

$$P_i(t+1) = \begin{cases} K \cdot P_i(t) + EP_p, & \text{if } x(t) \geq 0 \\ K \cdot P_i(t) + EP_N, & \text{otherwise} \end{cases} \quad (4.5)$$

K is a decay parameter that determines how much the perseveration value of each deck

is discounted in each trial. The tendency to persevere or switch is incremented each time and is updated by a gain impact parameter $EP_P(-Inf < EP_P < Inf)$ and a loss impact parameter $EP_N(-Inf < EP_N < Inf)$ based on whether the net outcome on the previous trial was a loss or a gain. Positive values for these parameters indicate stronger tendencies for decision makers to persevere with the choice of deck from the previous trial, whereas negative values indicate switching tendencies. The expected value and the perseverance term are then integrated into a single value signal:

$$V_i(t+1) = w \cdot E_i(t+1) + (1-w) \cdot P_i(t+1) \quad (4.6)$$

where $w(0 < w < 1)$ is a weight parameter that controls the weight given to the expected value in each trial. A greater value of w represents a greater weight given to the expected value. The values of $V_i(t+1)$ are then entered into the softmax function to calculate the probabilities of being chosen for each option.

The ORL model. The latest proposed ORL model assumes that people track the expected value ($EV(t)$) and win frequency ($EF(t)$) separately. Additionally, for positive and negative net outcomes, the decision makers update the expectancy of the chosen deck i with different learning rates.

$$EV_i(t+1) = \begin{cases} EV_i(t) + A_+ \cdot (x(t) - EV_i(t)), & \text{if } x(t) > 0 \\ EV_i(t) + A_- \cdot (x(t) - EV_i(t)), & \text{otherwise} \end{cases} \quad (4.7)$$

where A_+ ($0 < A_+ < 1$) and A_- ($0 < A_- < 1$) are the reward and punishment learning rates used to update the expected value of the chosen deck after rewards and punishment respectively. The updating process for win frequency ($EF(t)$) for the chosen option is as follows:

$$EF_i(t+1) = \begin{cases} EF_i(t) + A_+ \cdot (\text{sgn}(x(t)) - EF_i(t)), & \text{if } x(t) > 0 \\ EF_i(t) + A_- \cdot (\text{sgn}(x(t)) - EF_i(t)), & \text{otherwise} \end{cases} \quad (4.8)$$

in which A_+ and A_- are the same learning rates as those used to update the expected

value; $sgn(x(t))$ returns 1, 0, or -1 for positive, 0, or negative outcome values on trial t , respectively. The expected outcome frequencies for unchosen decks j' are also updated on each trial, in which the learning rates are also shared from the expected value learning rule:

$$EF_{j'}(t+1) = \begin{cases} EF_{j'}(t) + A_- \cdot \left(\frac{-sgn(x(t))}{C} - EF_{j'}(t)\right), & \text{if } x(t) > 0 \\ EF_{j'}(t) + A_+ \cdot \left(\frac{-sgn(x(t))}{C} - EF_{j'}(t)\right), & \text{otherwise} \end{cases} \quad (4.9)$$

C is the number of alternative choices for the chosen deck j , which is 3 in the case of the IGT. The ORL model also assumes that the decision makers have tendencies to stay or switch their choices regardless of the outcome in the last trial and this tendency can be captured by a perseverance weight ($PS_i(t)$).

$$PS_i(t+1) = \begin{cases} \frac{1}{K+1}, & \text{if } D(t) = i \\ \frac{PS_i(t)}{1+K}, & \text{otherwise} \end{cases} \quad (4.10)$$

where K is the decay parameter that controls how quickly the past deck selections are forgotten and it is determined by:

$$K = 3^{K'} - 1 \quad (4.11)$$

$K' \in [0, 5]$, so $K \in [0, 242]$. A single value signal for each deck is then produced via integrating the expected value, frequency, and perseverance into a linear function:

$$V_i(t+1) = EV_i(t+1) + EF_i(t+1) \cdot \beta_F + PS_i(t+1) \cdot \beta_P \quad (4.12)$$

$\beta_F(-Inf < \beta_F < Inf)$ and $\beta_P(-Inf < \beta_P < Inf)$ are two weight parameters that reflect the weight given to the outcome frequency and perseverance respectively relative to the expected value of each deck. Finally, the choice probability for each choice is determined by passing the expected values through the softmax function.

$$P[D(t+1) = i] = \frac{e^{V_i(t+1)}}{\sum_{k=1}^4 e^{V_k(t+1)}} \quad (4.13)$$

4.3 Results

Both individuals living with chronic pain ($N = 20$) and healthy controls ($N = 25$) completed the web-based BPI and the IGT at a location of their choice and at a time of their choice without supervision. The full demographic information is listed in Table 4.1.

Table 4.1: Demographic information for healthy controls and people living with chronic pain

	Healthy controls				Chronic pain				<i>P value</i>
N	25				20				0.067
Female	12				14				
	Mean	s.d.	Min	Max	Mean	s.d.	Min	Max	
Age	37.2	12.5	24	63	40.2	11.9	23	62	0.431
Pain duration (months)	1.2	1.3	0	5	78.6	76.3	6	264	<0.001
Pain severity	2.07	2.5	0	10	4.3	2.3	0	8	0.004
Interference REM	2	2.9	0	9.3	4.6	3.3	0	9	0.007
Interference (WASW)	1.9	2.6	0	8.3	4.2	2.5	0	8.5	0.004

4.3.1 Self-report Analysis

As expected, individuals with chronic pain demonstrated higher levels of pain severity ($t(43) = -3.06$; $P = 0.004$; Cohen's $d = -0.92$; 95% confidence interval (CI): $[-3.64, -0.75]$) and their daily activities were more influenced by pain ($t(43) = -3.05$; $P = 0.004$; $d = -0.92$; 95% CI: $[-4.1, -0.83]$), in comparison with healthy controls. Moreover, individuals living with chronic pain reported higher levels of sub-dimensional interference on REM ($t(43) = -2.86$; $P = 0.007$; $d = -0.86$; 95% CI: $[-4.50, -0.78]$) and WASW ($t(43) = -3.02$; $P = 0.004$; $d = -0.91$; 95% CI: $[-3.84, -0.77]$), respectively. Pain severity, as expected, strongly correlated with pain interference ($r(40) = 0.84$, 95% CI: $[0.73, 0.91]$, $P < 0.001$). The two sub-dimensional interferences were positively correlated ($r(43) = 0.88$, 95% CI: $[0.80, 0.93]$, $P < 0.001$) as well (Figure 4.2).




Figure 4.2: Pain severity plotted against pain interference (left); activity sub-dimension plotted against affective sub-dimension (right). Healthy controls ($N = 25$) are plotted in red and people living with chronic pain ($N = 20$) are plotted in blue. The r values were calculated between the paired pain measures for the whole sample ($*P < 0.5, **P < 0.01, ***P < 0.001$).




Figure 4.3: Pain severity plotted against total amount of gain by the end of the task. Healthy controls ($N = 25$) are plotted in red and people living with chronic pain ($N = 20$) are plotted in blue. The r value was calculated between the pain measure and the task performance measure for the whole sample ($*P < 0.5, **P < 0.01, ***P < 0.001$).

4.3.2 Standard Behavioral Data Analysis

The total amount of gain at the end of the IGT did not significantly differ between individuals with chronic pain (1997 ± 1187) and healthy controls (1756 ± 645 ; $t = 0.81$, $P = 0.42$). However, the pain severity was significantly correlated with the total gain ($r = -0.35$, 95% CI: $[-0.59, -0.07]$, $P = 0.01$, Figure 4.3). Figure 4.4 shows the proportion of choices from each deck as a function of the 5 blocks for the healthy and chronic pain groups separately and the proportion of choices from the good and bad decks. The choice pattern of the chronic pain group was qualitatively different (visual inspection of plots) from that of the healthy controls, although both groups demonstrated an avoidance of the bad Deck A but a preference to the other bad Deck B. Both Deck B and Deck D featured low-frequency losses and were generally chosen more often than Deck A and Deck C which featured high-frequency losses. Both healthy and chronic pain decision-makers selected more cards from good decks than bad decks at the beginning of the task. After learning whether each deck was good or bad in the second block, healthy controls continued to select more from good decks than bad decks. The choices of chronic pain decision-makers, however, seemed to fluctuate more across the advantageous and disadvantageous decks throughout the task. The final proportion of good deck selection of healthy decision-makers was higher than that of chronic pain decision-makers.

Figure 4.5 shows the learning scores across the five blocks of the IGT. A learning process was apparent for the healthy control group, in which the learning score progressively improved across the five blocks. Although the learning scores of individuals with chronic pain also showed an increasing trend, there is a clear dip in Block 4. To quantify the group differences, we applied the repeated measures ANOVA test in the form of a 5 (block) \times 2 (health status) form to the learning scores. To our surprise, neither block ($F(4, 172) = 2.37$, $P = 0.054$) nor group ($F(1, 43) = 0.31$, $P = 0.58$) factor had a significant impact on deck selection and there was no significant two-way interaction either ($F(4, 172) = 0.32$, $P = 0.87$). The results suggested that people living with chronic pain and healthy controls did not show significant deck preferences on the IGT and neither group developed a strong learning curve during the task.

4.3 Results

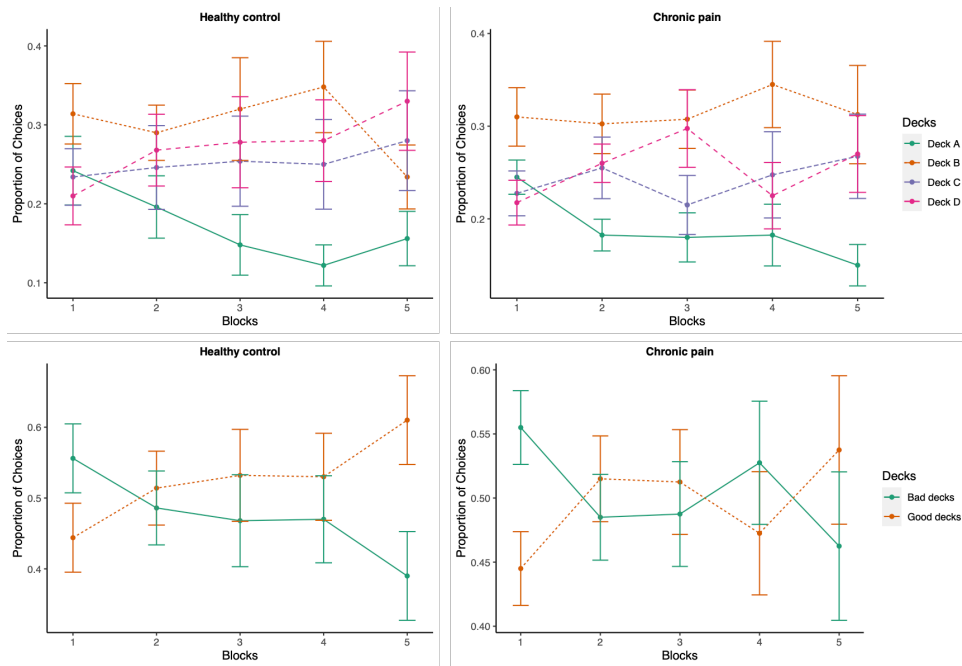


Figure 4.4: Mean proportion of choices from each deck within 5 blocks of both groups of decision makers (top two graphs); mean proportion of choices from good decks and bad decks of both groups of decision makers (the last two graphs). Each block contains 20 trials.




Figure 4.5: The learning IGT scores across the five different blocks of the IGT in healthy controls and people living with chronic pain.

4.3.3 Computational Modelling Analysis for the IGT

Even though the behavioral data statistics suggested that healthy and chronic pain decision makers did not show significant different deck preferences on the IGT, there might still be group differences in the cognitive processes underlying the decision choices. To investigate this possibility, we decomposed the IGT performance of the two groups using the cognitive modelling analysis introduced earlier. We first checked which model provided the best short-term prediction performance, as measured by the one-step-ahead PSIS-LOO. The smaller a model's PSIS-LOO score is, the better its model fit is. As shown in Table 4.2, the VPP model demonstrated the best overall model-fit relative to other models, followed by the ORL model. The ORL model ranked the best for the chronic pain group.

Table 4.2: Models and prior fits

Model	Pain $_{PSIS-LOO}$	Healthy $_{PSIS-LOO}$	Sum $_{PSIS-LOO}$
ORL	4593	4600	9194
VPP	4623	4544	9168
PVL_decay	4756	4867	9624
PVL_delta	5054	5622	10676

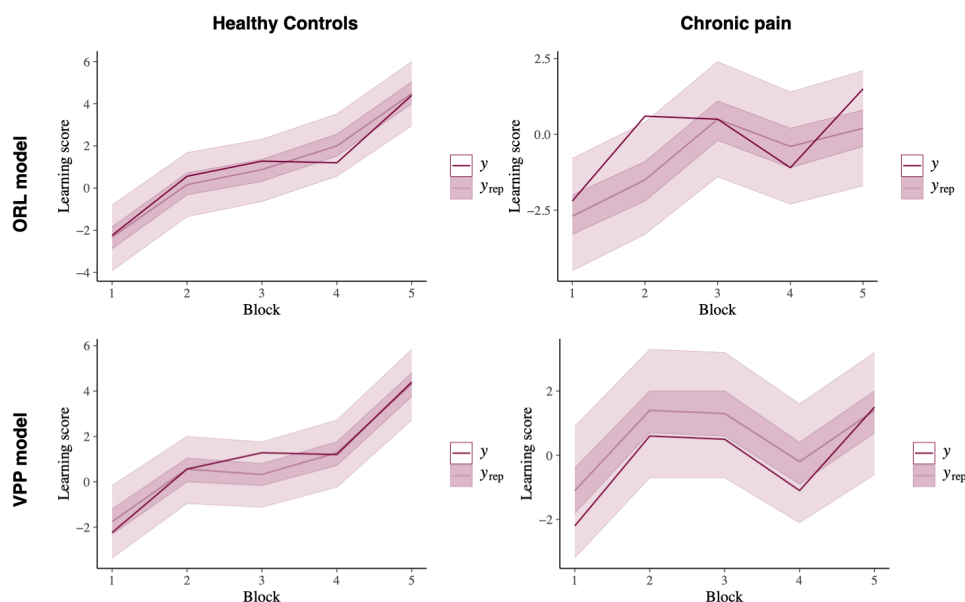


Figure 4.6: Model validation with PPC on the learning scores across five blocks for the ORL model and the VPP model. Model predictions plotted against actual learning scores. Shaded area depicts the 95% HDI of the posterior distribution.

To further validate the two best fitting models, PPC and parameter recovery analysis

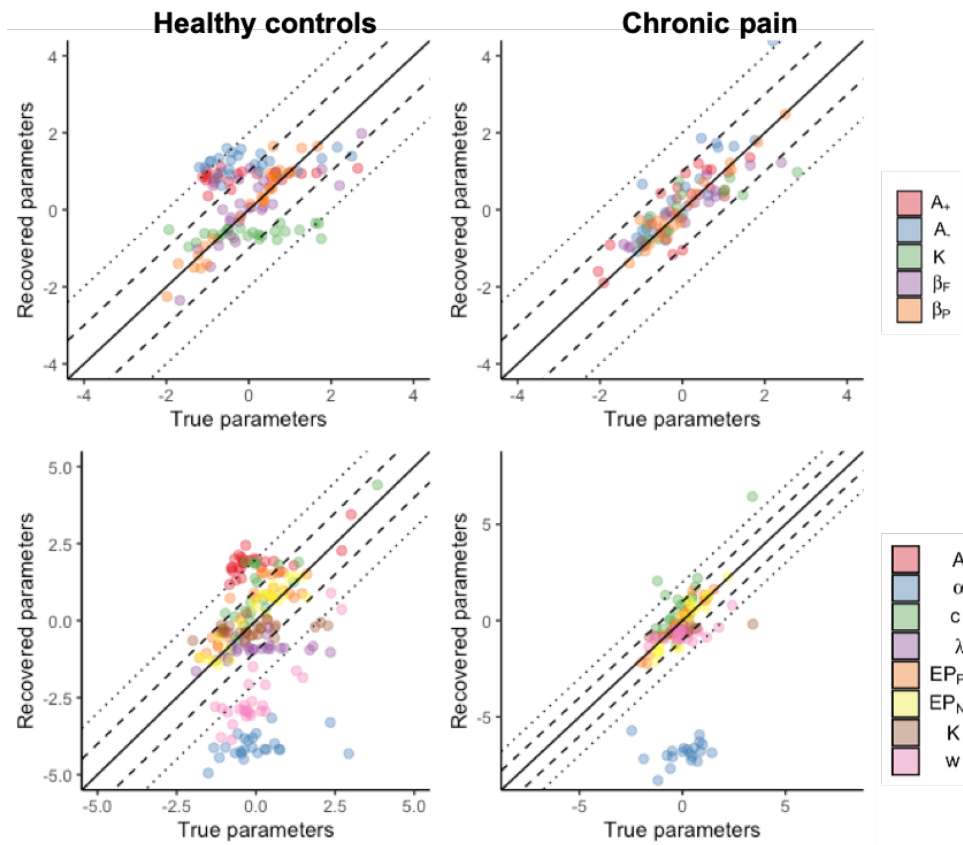


Figure 4.7: Parameter recovery of the ORL model (the top pannel) and the VPP model (the bottom pannel) for the chronic pain and the healthy participants. Both the means of the posterior distributions of the true and the recovered parameters were being standardized. Dashed and dotted lines reflect 1 and 2 standard deviations in the standardized space, respectively.

were conducted. In PPC analysis, Each model was used to generate synthetic choice data per trial and per participant for 12,000 times. The learning scores across five blocks was chosen as the model-independent measure to summarize both the real and the replicated choice data for comparing the predictive performance of the models. Figure 4.6 demonstrates the PPC results of the two models. The correlation between the replicated data and the real data for healthy control group is significant for both of the two models (ORL model, $r = 0.72$, $P = 0.003$; VPP model, $r = 0.68$, $P = 0.004$), while the correlation for chronic pain group is only significant for the VPP model (ORL model, $r = -0.47$, $P = 0.18$; VPP model, $r = 0.60$, $P = 0.006$). The true parameters (i.e., the parameter set used to simulate choices) used in parameter recovery were the means of the individual-level posterior distributions of each model (the ORL and the VPP) fitted to the chronic pain and healthy participants respectively. The recovered posterior means of each parameter were plotted against to the true parameters in a standardized space, in which both the true and the recovered posterior means of each parameter were z-scored by the mean and standard deviation of true parameters (Figure 4.7). This visualization enables us to observe the bias in recovered posterior means, where any values falling above or below the solid diagonal line indicate higher or lower recovered means in reference to the true parameters, respectively [135]. The recovered posterior means of the VPP model for the healthy group were systematically higher than the true parameters for the learning rate (A), and systematically lower for the outcome sensitivity parameter (α) and the weight parameter (w). The recovered means of the weight parameter (w) of the VPP model for the chronic pain group were also lower compared the true means. For the ORL model, the recovered posterior means were well-distributed around the true parameter means. The ORL model was selected to compare the two groups in the following analysis.

Figure 4.8 shows the posterior distributions of the group-level mean parameters of the ORL model fitted with two priors (one for each group) for healthy and chronic pain decision makers. The chronic pain group also showed strong evidence of increased reward learning rate (95% HDI from 0.22 to 0.55), punishment learning rate (95% HDI from 0.03 to 0.11) and the difference between the reward and punishment learning rate (95% HDI

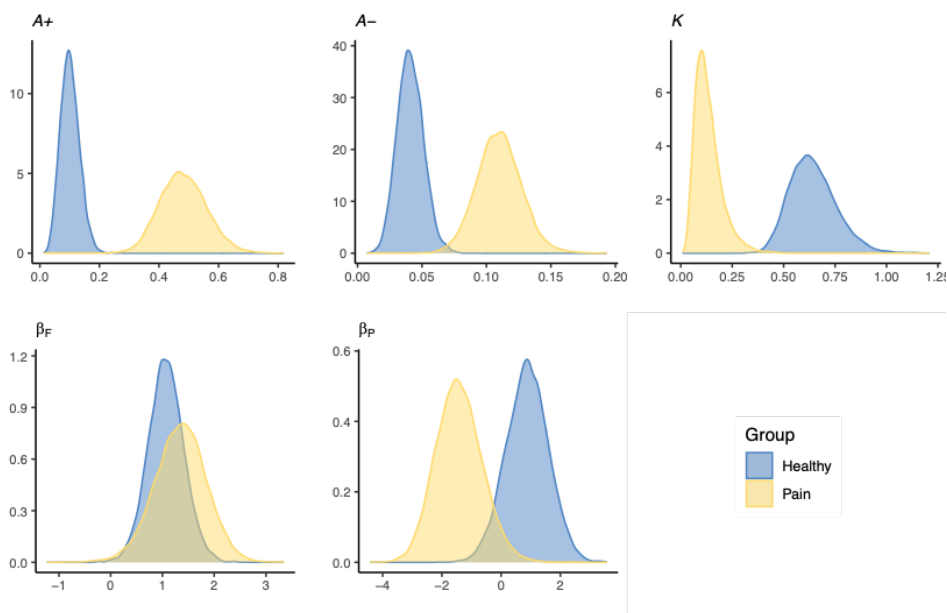


Figure 4.8: Groups-level ORL parameters across healthy controls and people living with chronic pain.

from 0.16 to 0.47), and decreased decay rate (95% HDI from -0.76 to -0.25), and outcome perseverance (95% HDI from -4.38 to -0.21) than healthy controls when fitting on the ORL model (the 95% HDI for the comparison across groups did not overlap zero, see Figure 4.9).

Extracting individual estimations for the ORL model provided evidence for the existence of positive correlations between pain severity the and reward learning rate parameter (see Figure 4.10; $r(43) = 0.41$, 95% CI: [0.12, 0.62], $\log BF_{10} = 2.18$, $P = 0.005$). Similar correlations were obtained between pain interference (including two sub-dimensional interference) and this model parameter. However, a negative correlation was observed between pain interference and the decay rate parameter ($r(43) = -0.29$, 95% IC: [-0.54, -0.00], $\log BF_{10} = 0.18$, $P = 0.05$), but this correlation with decay rate parameter did not apply to pain severity. Similar correlation was only identified between REM and decay rate parameter ($r(43) = -0.29$, 95% IC: [-0.54, -0.00], $\log BF_{10} = 0.18$, $P = 0.05$), but no supported correlation for the WASW ($r(43) = -0.27$, 95% IC: [-0.53, -0.02], $\log BF_{10} = -0.04$, $P = 0.07$) was identified.

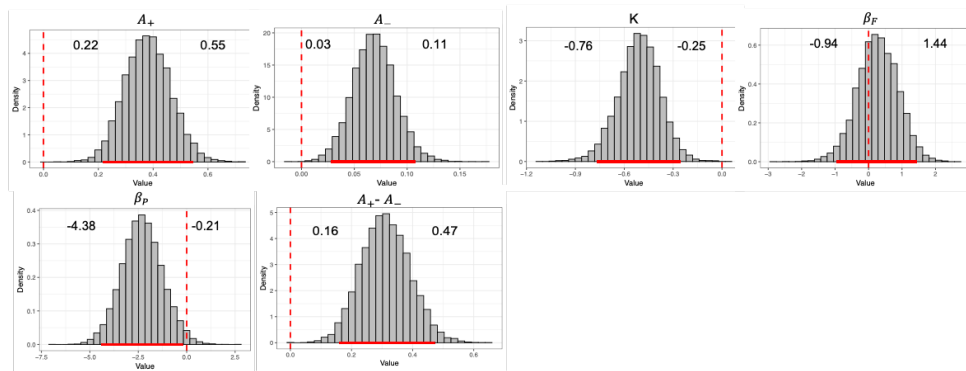


Figure 4.9: Differences in groups-level distributions for the ORL model between healthy and symptomatic groups. Solid red lines covers the 95% highest posterior density interval (HDI), and dashed red lines marks the 0 point. Values on the left and right sides of each graph are the lower and upper bounds of the 95% HDI of the comparison between the symptomatic and healthy control groups. If 0 point was included in the HDI, we consider there to be a non-significant difference between the two groups. We found a main effect of group on the reward learning rate (A_+), punishment learning rate (A_-), decay rate (K), perseverance weight (β_p), and the difference between the reward and punishment learning rates for the ORL model.

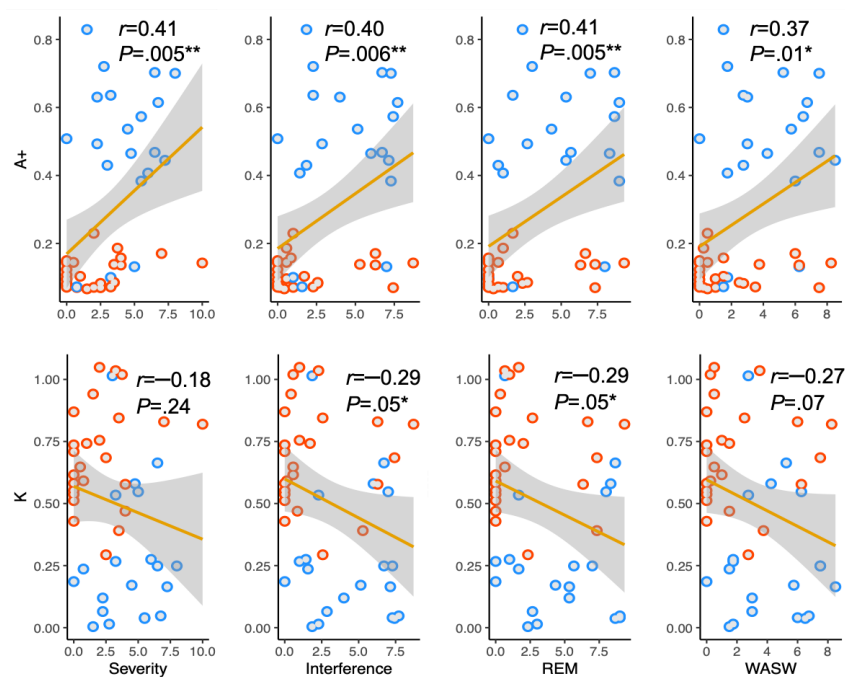


Figure 4.10: Estimated parameters of the ORL model plotted against four pain measures. Healthy controls are plotted in red and patients with chronic pain are plotted in blue. The r values were calculated between the model parameters and pain measures for the whole sample ($*P < 0.05$, $**P < 0.01$, $***P < 0.001$).

4.4 Discussion

In this study, we explored the differences in decision-making in individuals living with chronic pain against healthy controls in a field environment by collecting their behavioral responses to the web-based IGT. The main finding of our study was that people with chronic pain did not show significant differences in decision-making on the IGT based on the standard behavioral statistics. However, computational modelling analysis revealed that people with chronic pain had elevated learning rates both for rewards and punishments. Further results revealed chronic pain individuals were more dominated by rewards in the process of learning. Meanwhile, the symptomatic group demonstrated decreased decay rate and perseverance weight. We also explored the association among the self-report pain experiences, standard inferential statistics and cognitive parameters. The total amount of gain at the end of the task was negatively correlated with the degree of pain severity. Several cognitive parameters could also predict the pain severity and pain interference assessed by the self-report pain experiences.

Using standard inferential statistics, the total amount of gain obtained by the end of the task was not significantly lower in people living with chronic pain in comparison with healthy controls, indicating that there was no difference between the two groups in terms of the overall performance on the IGT. However, the total amount of gain was negatively correlated with the pain severity measure, suggesting that higher pain severity could be a factor that impairs performance on the IGT. Although people with chronic pain tended to select more bad decks relative to healthy controls as seen from the learning curves, the statistical analysis identified no significant differences between the two groups in terms of deck preference across blocks. There was lack of statistical evidence of significant learning across blocks for both of the groups either. These results are inconsistent with previous studies that identified significant group effects, in which patients won significantly less money and failed to adopt the advantageous decision-making strategy quickly learned by healthy controls [160], [168].

However, significant differences were identified in several cognitive components when comparing groups using the best-fitting model (ORL) in the computational modelling

analysis. The reward and punishment learning rates in the ORL model are used to update expectations after positive and negative outcomes, respectively. These two parameters account for the degree of the subject's sensitivity to losses and gains. Individuals with chronic pain demonstrated both elevated reward learning rates and punishment learning rates when fitting on the ORL model, suggesting they gave more weight to recent outcomes in the process of learning. It was suggested in [135] that comparing the difference between the reward and punishment learning rates for the two groups was more useful, although they were defined separately. The larger the differences between the two learning rates, the more the learning was dominated by either rewards or punishments. The significantly higher reward learning rates that caused larger differences between the two learning rates in individuals with chronic pain suggested that they were more sensitive to gains over losses relative to healthy controls. In other words, chronic pain individuals appeared to be more driven by rewards, which could be a possible reason that made them choose more cards from Deck B (as seen in Figure 4.5), the bad deck with a higher reward magnitude but also a higher punishment magnitude.

It can be seen from Equation (4.12), that the outcome frequency weight and perseverance weight parameter in the ORL model collectively influence the total value of the expected value of each deck. Values for the outcome frequency less than or greater than 0 indicate that decision makers prefer decks with low or high win frequency, respectively. The outcome frequency for the two groups were both larger than 0 and did not show significant differences, suggesting both of them prefer decks with high win frequency, i.e. Deck B and Deck D, which can be reflected in the learning curves plotted in Figure 4.5 where healthy individuals ended up selecting Deck D most and patients ended up selecting Deck B most. Values for the perseverance weight parameter less than or greater than 0 indicate that decision-makers prefer to switch or stay with their recent chosen decks. The mean value of this parameter for the healthy controls was larger than 0, whereas this figure for the patients was smaller than 0. Meanwhile, there was a significant difference between the two groups. It means people with chronic pain were less persistent in their previous choices during the task. In this sense, our findings were in agreement with a previous study [161]

where a simple heuristic model was proposed to discriminate between patients and healthy controls on IGT performance by tuning the degree of randomness and importance given to losses and gains. Patients with chronic pain in their study demonstrated significantly less persistent behavior, which was characterized by giving more emphasis to gains than losses and increasing decision randomness. Contrary to our expectation, however, chronic pain decision makers presented lower decay rates, suggesting that they could remember longer histories of their own deck selections relative to healthy controls. However, substantial existing studies have reported impaired memory functions in patients with chronic pain [169], [170] and memory complaint was one of the most common complaints in patients with chronic pain with cognitive deficits [171].

When analyzing the correlations between the cognitive parameters and self-report pain experiences, we found that higher reward learning rates in the ORL model could significantly predict higher self-reported pain severity and pain interference. Lower decay rates were only associated with higher self-reported pain interference, especially with the affective subdimension of interference. Given the correlations observed between the cognitive parameters and pain experiences, the cognitive task might be an important tool to take into consideration in evaluating individuals at risk of developing chronic pain conditions.

4.5 Conclusion

In summary, by recruiting participants online and administering a lab-in-the-field experiment, we found that chronic pain subjects displayed increased reward sensitivity and reduced choice persistency to their previous choices relative to healthy controls and revealed credible associations between the computational parameters and pain severity and pain interference. Compared to conventional statistical analysis of behavioral performance, computational modelling analysis revealed much more evidence of distinct differences in decision-making between individuals with chronic pain and healthy controls in the noisier lab-in-the-field environment. The successful implementation and significant findings of this study provide experimental evidence for the feasibility of measuring and comparing

the cognitive performance of people with chronic pain and healthy controls in their daily environments instead of in laboratories settings. Moving the behavioral paradigms outside of the lab to the field is a crucial step to improve ecological valid of decision-making studies, allowing us to observe how people make decisions when they are not under any control or supervision. It has implications for researchers who are interested to explore whether the cognitive alterations of chronic pain individuals may predict their functioning in everyday life. More broadly, the benefits of shifting from the lab to the field may extend to clinical outcomes, facilitating early detection of cognitive impairments and personalized treatment, although before that the equivalence of the paradigms in the lab and field requires more careful empirical testing and validation both in lab and outside the lab in various settings.

Chapter 5

Impaired Adaptation of Learning to Contingency Volatility is Attributed to the Depression and Anxiety Symptoms of Parkinson's Disease: A Lab-in-the-field Study

This chapter aims to further validate the feasibility of evaluating decision-making behavior in the real-world settings by applying this methodology to patients with Parkinson's Disease (PD), a population who may require regular care and supports from others for daily functions. Theoretical and empirical evidence has suggested that patients with PD showed various deficits in feedback learning due to the depletion of dopaminergic neurons and excessive dopaminergic medication, which are involved in the process of prediction error coding in the striatum that is critical for the learning process. Research on the impact of non-motor symptoms of PD is limited, regardless of the fact that non-motor symptoms are in fact significant burdens on the patients and their caregivers. Particularly, two of the most common neuropsychiatric symptoms, i.e. depression and anxiety which are difficult to recognize and are under-treated in clinical settings, have been associated

with biased reward/punishment learning behavior. In this study, we examined how patients with PD performed in terms of adaptation of learning rates to contingency volatility on the web-based ALT and if their performance was attributed to the severity of the non-motor symptoms of PD. Consistent with the findings in the general population, patients with PD with higher anxiety and depression levels were less able to adjust their learning rates on the basis of whether action-outcome associations were stable or volatile in a simulated aversive environment, but this deficit was not correlated with the severity of PD or other non-motor symptoms. The depression- and anxiety-specific deficits of adaptation learning to the contingency represented by the computational parameters provide more insight into the nature of learning disturbances in patients with PD. The computational predictors detected in this study may facilitate early detection and contribute to future individualized treatment of the two symptoms comorbid to PD.

5.1 Introduction

Parkinson's disease (PD) is a chronic, progressive neurodegenerative disorder due to the loss of dopaminergic neurons within the substantia nigra located in the midbrain [172]. It is the second most common neurodegenerative disease, with approximately 2% of the world's population aged 65 years or older afflicted by it [173]. While classically defined as a movement disorder that is characterized by degeneration of mobility and muscle control, the neurotransmitter deficits in PD can also lead to various non-motor symptoms that have great impacts on motivational drive and reward processing behavior, which themselves are fundamental to the learning of new responses and decision-making.

As introduced in **Chapter 2**, a broad body of theoretical and empirical evidence has suggested that phasic activity of the midbrain dopamine system is critical to incremental experience-based learning and decision-making [78], [87], [88], [174], [175]. The dopamine neurons are known to encode a reward-related signal called reward prediction error that indicates a discrepancy between obtained and expected reward values. This prediction error, in turn, is thought to play a key role in rewarded learning and has gained widespread use in temporal difference models of learning that are driven by reinforcing

rewards [58]. Electrophysiological studies in primates showed that reward elicits phasic dopamine increases, whereas aversive feedback leads to phasic dopamine decreases in midbrain dopamine neurons. In other words, the reduction of phasic dopamine responses, due to a decrease or an excessive increase of dopamine level, could negatively affect prediction error and, consequently, impair reinforcement learning and decision-making. Thus, the degeneration of dopamine cells in PD, which not only reduces the tonic level of dopaminergic activity but also impairs phasic dopaminergic activity, has been associated with impaired reward-based learning [176]–[178]. Although the administration of dopaminergic medication ameliorates the motor deficits in PD, according to the overdose hypothesis, the excessive supplement of dopamine may prevent dopamine dips that are critical for error detection, thus leading to a specific impairment in learning from negative feedback [177].

Multiple controlled medication withdrawal studies in PD have tested the ON/OFF state of medication on various types of learning and decision-making paradigms [179], [180]. ON state is defined as when patients are mobile with a lessening of the parkinsonian symptoms, whereas OFF condition is when the antiparkinsonian medications are wearing off and parkinsonian symptoms reappear, i.e. patients show classical tremor or akinetorigid presentations [181]. Compatible with the overdose hypothesis, a predominant finding was that dopamine medication in PD enhanced reward-based learning that involves learning to select responses that lead to rewards, while reduced punishment-based learning that leads to negative outcomes in probabilistic learning paradigms [182]–[186]. Additionally, patients with PD have been found to have difficulty learning the action-outcome contingency reversals, particularly from the negative outcomes in reversal learning paradigms, and to be more likely to commit perseveration errors while making decisions in a medicated state. These effects were compelling, both theoretically and empirically, but several recent studies have failed to replicate them [187], [188]. While at first puzzling, this variability in the effects of dopaminergic medication is perhaps not so surprising. It concurs with extensive evidence from the field of clinical psychology that the neuropsychiatric non-motor symptoms, e.g. depression and anxiety, which are highly prevalent in PD, have been

independently associated with altered reward processing among the general population without PD [189]–[192]. However, knowledge about the influence of these psychiatric factors on reinforcement learning in the PD population is limited.

The impact of the neuropsychiatric symptoms like depression and anxiety extends far beyond mood symptoms in PD: earlier initiation of dopaminergic therapy, greater functional disability, faster physical and cognitive deterioration, increased mortality, poorer quality of life, and increased caregiver distress [193], [194]. However, in clinical settings, depressive and anxious disturbances are often under-recognized and, even when identified, frequently under-treated [195]. For example, one of the reasons for the difficulty of detection of depression and anxiety among patients with PD is that many somatic symptoms of these two syndromes are also evident in patients with PD without depression and anxiety disorder [196]. Pathophysiological changes involved in the pathogenesis of PD are also suspected to play a prominent role in depression in the general population. A prevailing model for the development of depression in PD proposes that the degeneration of mesocortical and mesolimbic dopaminergic neurons causes orbitofrontal dysfunction, which disrupts serotonergic neurons in the dorsal raphe and leads to dysfunction of depression-related orbitofrontal-basal ganglia-thalamic circuits [197]. In fact, the interaction between depression heralding PD and incidental depression after the diagnosis of PD is complex and a unified pathophysiological model of depression in PD is lacking. Moreover, treatment of non-motor symptoms is often difficult and may either improve or worsen under dopaminergic replacement therapy used to alleviate motor symptoms. In order to develop and improve treatments for non-motor symptoms, it is necessary to enhance our understanding of the mechanisms underlying non-motor symptoms like anxiety and depression.

In this study, we ask how the non-motor symptoms, particularly depression and anxiety symptoms, may contribute to the learning deficits in PD. We aim to identify the correlation by conducting a lab-in-the-field experiment using the web-based ALT, a task that has shown good reliability and robustness for examining the relationship between anxiety and depression and adaptation of decision-making to contingency volatility in general population without PD [198]. The web-based nature of the experiment allows the patients

to provide assessment data using their internet-enabled hardware without an extra clinical interview, which is particularly valuable for patients with PD who may have difficulties in moving and traveling. Computational-based analysis allows us to distinguish between different potential causes for impairment in dealing with the uncertainties involved in the ALT. Two sources of uncertainties have been formalized in the literature of computational modelling. One source of uncertainty is the noise in the relationship between actions and outcomes, while the second source of uncertainty is the changes in action-outcome contingencies in the underlying structure. Subjects are supposed to adopt higher learning rates when action-outcome contingencies are volatile (i.e. the second-order uncertainty is high) than when contingencies are stable (i.e. the second-order uncertainty is low). Failure to adapt learning rates correctly to the source of uncertainty present in a given situation may result in inaccurate predictions and suboptimal decisions. Depression and anxiety in the general population have been associated with impoverished adjustment of learning to volatility. Our question is whether this association can be generalized to the depressive and anxiety symptoms of the PD population and whether PD itself leads to such deficits in the adaptation of learning to contingency volatility. If such deficits are unique predictors for depressive and anxiety symptoms, the altered computational phenotypes may facilitate early detection and the development of the optimal care that is needed for patients with these disturbances in PD.

5.2 Methods

5.2.1 Recruitment

PD subjects were recruited by the practitioners involved in the General Neurology and Movement Disorder Program, Foothills Medical Center, Alberta Health Services in the Calgary area, in the state of Alberta, Canada. Men and women patients in any phase of the disease, who accepted to participate in the study and signed the informed consent form were included. All patients were under dopaminergic medication. Inclusion criteria include the diagnosis of idiopathic Parkinson's disease according to the UK Brain Bank diagnostic

criteria [199]. Exclusion criteria include patients with a diagnosis of dementia (Montreal Cognitive Assessment < 20), participants who cannot engage with the decision-making tool and participants who are unable to speak or understand English. This is a correlational study. No healthy controls were involved for comparison since patients with PD in this study were required to do a simpler version of the ALT to reduce their burden and therefore a control group comparison was not possible. The study was approved by the human research ethics committee of University of Calgary. Participants Subjects were advised both verbally and in the written plain language statement and informed consent when they were recruited in the study.

5.2.2 Study Design

The primary outcome will be the relationship between the ability to adjust the learning rate to contingency volatility and the presence and severity of non-motor symptoms of PD, including depression, anxiety, apathy, pain and fatigue. Symptom evaluation was administered during the clinical interview using validated instruments because appropriate training in the administration and scoring of the scale is necessary to obtain reliable scores for some of the instruments [200]. Additional information collected included demographic information (age, sex and gender, educational level, cultural background, occupation), medical history and a list of medications. The entire duration of the visit was 60-90 minutes. The adaptive learning ability was assessed with the web-based ALT, which means the patients could choose to complete the task at a time and location of their own choices without supervision from the practitioners.

- **The Parkinson Disease Questionnaire (PDQ)** is a PD-specific health status and quality of life questionnaire and is the most frequently used disease-specific health status measure in clinical trials. It comprises 39 items assessing how often people with Parkinson's experience difficulties across 8 dimensions including mobility (10 questions), activities of daily living (6 questions), emotional well-being (6 questions), stigma (4 questions), social support (3 questions), cognition (4 questions), communication (3 questions), and bodily discomfort (3 questions) [201]. Patients are

requested to answer the questions based on how often (from never to always) they have experienced the problem defined by each item during the past month because of their PD.

- **MDS-Unified Parkinson’s Disease Rating Scale (MDS-UPDRS)** is the most used tool for rating PD. It has four parts: Part I (non-motor experiences of daily life), Part II (motor experiences of daily life), Part III (motor exploration) and Part IV (motor complications).
- **Hamilton Depression Rating Scale (HDRS)** is the most widely used and accepted interviewer-rated measure for evaluating the severity of the illness in patients suffering from depression without somatic comorbidity. The discriminant validity, sensitivity, specificity, and test-retest reliability of HDRS are very high. The original version contains 17 items pertaining to symptoms of depression experienced over the past week. Cutoff scores of 9/10 [202] and 11/12 [203] to screen for depression in PD, and 15/16 [202] and 13/14 [203] to diagnose major depressive disorder have been suggested. Using these cutoffs, sensitivity, specificity, and positive and negative predictive values for a DSM-IV diagnosis of major depressive disorder in PD have been found to be acceptable. It has been demonstrated to be sensitive to change in patients with PD [204], [205] and to correlate with biological markers of depression in PD.
- **Hamilton Scale for Anxiety (HAM-A)** is one of the first rating scales developed to measure the severity of anxiety symptoms and is still widely used today in both clinical and research settings. The scale consists of 14 items, each defined by symptoms, and measures both psychic anxiety (mental agitation and psychological distress) and somatic anxiety (physical complaints related to anxiety). Each item is scored on a scale of 0 (not present) to 4 (severe), with a total score range of 0 – 56, where < 17 indicates mild severity, 18 – 24 mild to moderate severity and 25 – 30 moderate to severe severity. It has been demonstrated good psychometric properties and internal consistency in patients with PD with general anxiety disorder [206].

- **King's Parkinson's Disease Pain Scale (KPPS)** is a reliable and valid scale for grade rating of various types of pain in PD. It is a scale developed in 2015 by Chaudhuri et al. [207], to assess pain specifically in the PD population. This rater-based 14-item scale measures the severity and frequency of pain in 7 various domains. Each item is scored by severity from 0 (none) to 3 (very severe) multiplied by frequency from 0 (never) to 4 (all the time) resulting in a subscore of 0 to 12, the sum of which gives the total score with a theoretical range from 0 to 168. Acceptable psychometric properties in PD have been reported for this scale.
- **The Starkstein Apathy Scale (SAS)** is the most used psychometric tool for assessing apathy in patients with PD. The SAS consists of 14 items (total score range 0 – 42; higher scores indicate more severe apathy) phrased as questions to be answered on a four-point Likert scale (intensity scores: not at all, slightly, some, a lot). In the first eight questions, the "not at all" answer corresponds to severe apathy; for questions 9 – 14, the intensity scores are reversed, such that the "a lot" answer corresponds to more severe apathy.
- **The web-based ALT** In order to reduce the burden for the patients with PD considering about their health conditions, the number of trials of the task was reduced to 120 trials, including 60 stable trials and 60 volatile trials. Other features remained unchanged as described in **Chapter 3**

5.2.3 Model-agnostic Behavioral Analysis

For each participant, we measured four basic trial-by-trial model-agnostic behavioral metrics: (1) the probability of choosing the favorable option with a lower probability of stealing (correct choice), (2) the probability of choosing the option with a smaller stealing magnitude (smaller choice), (3) the proportion of trials in which the participant selected a different choice after suffering from loss in previous trials (loss-shift), and (4) the proportion of trials in which the participant stayed with the same choice when there was no loss in previous trials (win-stay). Each of the metrics was calculated separately for

stable and volatile blocks and was submitted to repeated-measures ANOVA analysis to assess the impact of volatility of the environment on behavioral-level performance. Scores of non-motor symptoms in patients with PD were added as covariates in the analysis. Regression analysis was conducted to examine the correlation between various non-motor symptoms and behavioral summaries of the decision-making performance on the task. The relationship of the model-agnostic measures with the model-derived parameters (obtained in the next step of analysis) was also examined.

5.2.4 Computational Modelling Analysis

The models considered in this study were informed by prior work. Each of the models includes a parameter to allow for individual differences in the weighting of outcome probabilities against outcome magnitudes, an inverse temperature parameter to allow for differences in the degree to which choices are in line with this weighted combination of probability and magnitude, and a learning rate parameter that captured the extent to which participants update probability estimates given the unexpected reward loss in previous trial's outcome. In total, five alternative models were assessed and compared.

One important cognitive process in a probabilistic decision-making paradigm is to evaluate the SVs, integrating the probability and potential loss magnitude information. The first model (model 1) [104] assumes that participants combine outcome probability and outcome magnitude multiplicatively during decision-making, while the second model (model 2) [208] supposes that participants combine them additively. Model 1 supposes the probability P_t that a bad outcome (loss of coins) would result from choosing the blue leprechaun rather than the red leprechaun, is updated on a trial-by-trial basis using the RW rule.

$$P_t = P_t + \alpha(O_{t-1} - P_{t-1}) \quad (5.1)$$

in which the learning rate $\alpha \in (0, 1)$ determines how much weight the decision maker gives to the recent outcomes when updating their expected probability. The outcome O_{t-1} is coded as 1 if the blue leprechaun was chosen and produces a bad result or if the leprechaun

red was chosen and followed by a good result. O_{t-1} is coded as 0 for the opposite situation in the task. The initial outcome probability P_t was set as 0.5.

The outcome probability estimate is then adjusted to P'_t using a risk preference parameter ($\gamma \in (0, 10)$) to capture the relative importance of the magnitude of losses versus outcome probability. If $\gamma < 1$, it means that the experiment participant places greater weight on the magnitudes of losses, while if $\gamma > 1$, it means that the participant places greater weight on outcome probability when performing a choice.

$$P'_t = \min[\max[(\gamma(P_t - 0.5) + 0.5), 0], 1] \quad (5.2)$$

The expected value for each leprechaun is then calculated by multiplying the adjusted outcome probability and loss magnitude, separately, before taking the difference in expected value between the two leprechauns.

$$v_t = P'_t M_t^{blue} - (1 - P'_t) M_t^{red} \quad (5.3)$$

Finally, the action probabilities are generated using a softmax function with an inverse temperature parameter β , which controls the degree to which the expected values are used in choosing the leprechauns.

$$P(C = blue) = \frac{1}{1 + \exp(\beta v_t)} \quad (5.4)$$

Model 2 uses the same assumption as model 1 in terms of updating the outcome probability (the RW rule). However, in contrast with model 1, model 2 assumes the decision-makers combine outcome probability and outcome magnitude additively using a mixture weight parameter (λ).

$$v_t = \lambda [P_t - (1 - P_t)] + (1 - \lambda) [M_t^{blue} - M_t^{red}] \quad (5.5)$$

Model 3 is the same as model 2 except it assumes participants might treat differences in outcome magnitudes non-linearly (i.e. that they make decisions on the basis of subjective

rather than objective outcomes magnitudes). Thus, the difference in outcome magnitude is nonlinearly scaled with a scaling parameter ($r \in (0.1, 10)$) to capture any potential bias that the participants have towards treating differences in outcome magnitudes in model 3.

$$v_t = \lambda [P_t - (1 - P_t)] + (1 - \lambda) [M_t^{blue} - M_t^{red}]^r \quad (5.6)$$

The difference between model 4 and model 2 is that model 4 incorporates a choice kernel k_t . This choice kernel acts like an average window moving forward as the trials proceed. It is updated by an update rate parameter ($\eta \in (0, 1)$), which can be used to determine the number of recent choices contained in the value of the choice kernel on the current trial.

$$k_t = k_{t-1} + \eta (C_{t-1} - k_{t-1}) \quad (5.7)$$

The impact of the choice kernel on choice was determined by an additional inverse temperature parameter (ω_k) in a softmax function:

$$P(C = blue) = \frac{1}{1 + \exp(-(\omega v_t + \omega_k [k_t - (1 - k_t)]))} \quad (5.8)$$

Model 5 incorporates both the subjective scaling of the outcome magnitudes in model 3 and the choice kernel components in model 4 based on model 2.

Decomposing Parameters into Parameter Components

In order to capture the differences in choice behavior associated with the experimental conditions, i.e. the block type (volatile vs. stable) and relative outcome value (good vs. bad), the core parameters were decomposed into sub-components. Specifically, the learning rate parameter was divided into a baseline learning rate (α_{base}), a difference in learning rates between the volatile and stable blocks (α_{vs}), a difference in learning rates between good outcomes and bad outcomes (α_{gb}), and the two-way interaction of those differences:

$$\alpha = \text{logit}(\alpha_{base} + \alpha_{vs}\chi_{vs} + \alpha_{gb}\chi_{gb}) + \alpha_{vs \times gb}\chi_{(vs \times gb)} \quad (5.9)$$

in which a logistic transform was applied to constrain the overall learning rate α to be between (0, 1). This transform was also applied to other parameters that were constrained to be between 0 and 1. Other parameters involved in the models were not decomposed in order to avoid overfitting problems due to the introduction of too many parameters. Additionally, there is evidence from previous studies that the task volatility manipulation and relative outcome valence had no effect on them.

Hierarchical Bayesian Estimation of Parameters

A Hierarchical Bayesian procedure was used to estimate distributions over model parameters at both individual- and population-level for each alternative model. Data from all participants was used to fit each model. Specifically, each parameter was assigned an independent population-level distribution that was shared across participants. The standard deviation for the population-level distribution was estimated separately for each parameter. Posterior distributions were estimated using Hamiltonian Monte Carlo with an NoU-Turn Sampler (HMC with NUTS) as implemented in Stan [125] via its RStan interface. The Gelman-Rubin index \hat{R} (Rhat) [209] was used to assess the convergence of the MCMC samples. \hat{R} values close to 1.00 indicate that the MCMC chains have converged to stationary target distributions. There were no parameters with R values greater than 1.1 (most were below 1.01) in the process of sampling. Four chains were run with 1000 warming up and 4000 samples each.

Estimation of parameters using both group and individual level information may introduce dependencies among the individual level estimates, in that using participants' parameters on an individual basis to compare against outside measures (e.g., non-motor symptom severity) can be biased. Therefore, to examine relationships between parameters and variables of interest, the effects were estimated within the model by introducing another parameter to index the effect of the variable of interest. To do so, the non-motor

symptoms were taken as covariates and entered into a regression during model estimation to predict the mean of a parameter. It is notable that the non-motor symptoms were not fed to the regression model together in the same model because they are highly correlated. For example, the following analysis determines the effect of depression on the parameter that represents the difference in learning rates between the volatile and stable blocks (α_{vs}):

$$\alpha_{vs} \sim \text{Normal}(\mu + \beta_d X_d, \sigma) \quad (5.10)$$

Parameters were given a non-centred parameterization to increase the estimation performance by specifying mean, scale, and error distributions for each parameter. The group-level means were assumed to be normally distributed and the group-level scale parameters were specified as half-Cauchy distributed. Error distributions, which were estimated for each subject, were given a normal prior with mean = 0 and standard deviation = 1. Similarly, effects of covariates (e.g. β_d) were given a normal prior with mean = 0 and standard deviation = 1.

Model Comparison

We compared the fits of different models using PSIS-LOO, which is a popular method with which to estimate the out-of-sample prediction accuracy as it is less computationally expensive than evaluating exact leave-one-out or k-fold cross-validation accuracy. If the difference of the PSIS-LOO between the two models was larger than four times of the standard error of this difference, we consider there was a significant difference between the two models given that the five models were constructed through the inclusion of additional model parameters. PPC was also conducted to validate the absolute performance of the models. Each model generated synthetic choice data per trial and per participant for 12,000 times. The probability of correct choice was chosen as the model-independent measure to summarize both the real and the replicated choice data at three levels for comparing the predictive performance of the models. The correlations of the model-independent measure between the replicated data and the real data were calculated at trial level and subject level. The Bayesian p-value was calculated at the overall level. The winning model was

used to make inferences about the relationships between task performance and non-motor symptoms. Each model was estimated using the same hierarchical Bayesian procedure.

5.3 Results

5.3.1 Demographic and Clinical Data

Participants included 38 patients with a diagnosis of PD (20 males). The mean age was 62.70(\pm 9.63) years and the mean disease duration of the participants was 5.6 (\pm 3.95) years (see Table 5.1 for listed details of the demographic and clinical data). The UPDRS part III (the physical exam) could not be administered because in-person visits were not allowed during the pandemic. The scores on the other subscales of UPDRS and the Levodopa Equivalent Dose (LED) are reported here as an indicator of the severity of PD. The severity of depression symptom was significantly correlated with the severity of anxiety symptom ($r = 0.73, p < 0.001$). Both of them were significantly correlated with the severity of pain respectively (depression: $r = 0.57, p < 0.001$; anxiety: $r = 0.50, p = 0.002$). The apathy symptom is weakly correlated with the other three symptoms. All of the non-motor symptoms listed here are positively associated with the severity of PD as represented by the total score of UPDRS (depression: $r = 0.66, p < 0.001$; anxiety: $r = 0.78, p < 0.001$; pain: $r = 0.61, p < 0.001$; apathy: $r = 0.45, p = 0.005$) instead of the LED.

Table 5.1: Basic demographic and clinical details of the participants

Number of participants (Total N)	38
Female (N)	20
Age (mean \pm sd)	62.7 \pm 9.63
Disease duration (mean \pm sd)	5.61 \pm 3.95
LED (mg/day mean \pm sd)	757 \pm 519
PDQ (mean \pm sd)	37.29 \pm 20.89
UPDRSI (mean \pm sd)	12.66 \pm 6.32
UPDRSII (mean \pm sd)	12.05 \pm 6.57
UPDRSIV (mean \pm sd)	6.97 \pm 4.12
HDRS (mean \pm sd)	9.08 \pm 5.35
HAM-A (mean \pm sd)	14.26 \pm 7.50
KPPS (mean \pm sd)	23.45.29 \pm 18.77
SAS(mean \pm sd)	12.81 \pm 6.12




Figure 5.1: The proportion of trials selecting the option with a lower probability of yielding losses and with smaller loss magnitudes and the proportion of trials shifting when receiving or not receiving losses respectively. Each dot represents a participant and error bars represent 1SEM.

5.3.2 Model-agnostic Analysis

We examined several aspects of the choice data to investigate the learning patterns of patients with PD that are independent of the formal computational models (see Figure 5.1). Participants made significantly fewer correct choices in the volatile block compared to that in the stable block ($F = 39.55, P < 0.001$), while the preference for choosing the smaller options was not significantly changed transferring from stable to volatile block ($F = 0.05, P = 0.82$). Participants were more likely to shift to the other option in the volatile block after being stolen by the leprechaun ($F = 16.65, P < 0.001$), while the tendency of shifting was stable across blocks when there was no loss caused ($F = 2.96, p = 0.09$), indicating that participant behavior was more sensitive to bad outcomes in the volatile than the stable environment, as an agent with a higher learning rate would be.

The relationship of the model-agnostic measures of behavior with the PD and the non-motor symptom severity was also explored in this section. The difference of the proportion of choosing the correct choice ($\text{diff}_{\text{correct}}$) was slightly negatively correlated with all of the three non-motor symptoms (depression: $r = -0.38, p = 0.019$; anxiety:

$r = -0.36, p = 0.028$; apathy: $r = -0.4, p = 0.013$). The difference of the proportion of loss-shift measure ($\text{diff}_{\text{loss-shift}}$) was non-significantly related to depression and apathy scale, but slightly negatively correlated with anxiety level (depression: $r = -0.3, p = 0.065$; anxiety: $r = -0.37, p = 0.023$; apathy: $r = -0.21, p = 0.2$). The proportion of choosing the smaller options (*smaller*) throughout the task was not correlated with any of the non-motor symptoms. Particularly, none of the behavioral metrics was predicted by the severity of PD.

5.3.3 Computational Modelling Analysis

Model Comparison

We fitted five models to the data. Model 1 and model 2 were compared to investigate how participants integrated outcome probability and outcome magnitude (additive or multiplicative) while calculating the SVs. We observed that the additive model fitted participants' choice behavior better (model 2 PSIS-LOO = 4,523 vs. model 1 PSIS-LOO = 4672; difference in PSIS-LOO = -149; standard error of difference = 56). This finding is consistent with findings from multiple studies where additive models have also been reported to fit choice behaviour better than expected value models [198], [210]. Model comparison between model 2 and model 3 revealed that allowing subjective weighting of magnitude difference did not significantly enhance model fit (model 3 PSIS-LOO = 4521 vs. model 2 = 4523; difference in PSIS-LOO = -2.8, standard error of difference = 9.6). The addition of a choice kernel that captures participants' predisposition to repeating prior choices did not improve the model fit either based on the comparison results between model 2 and model 4 (model 4 PSIS-LOO = 4526 vs. model 2 = 4523; difference in PSIS-LOO = 2.8; standard error of difference = 6.4). Incorporating both subjective effects of weighting magnitude difference and choice kernel in model 5 did not improve the model fit either (model 5 PSIS-LOO = 4517 vs. model 2 = 4523; difference in PSIS-LOO = -6.2; standard error of difference = 10.6). The results of the PPC were consistent with the results of the LOOCV analysis (Figure 5.2 shows the PPC results plotted at three levels for model 1 and model 2.). The correlation coefficient between the replicated data and the real data

for model 2 was significantly improved compared to that for model 1 no matter at trial level or individual level (model 2, $r_{trial} = 0.85$, $P < 0.001$, $r_{subject} = 0.95$, $P < 0.001$, Bayesian $P_{overall} = 0.04$; model 1, $r_{trial} = 0.83$, $P < 0.001$, $r_{subject} = 0.90$, $P < 0.001$, Bayesian $P_{overall} = 0$), while the correlation coefficient for model 3, model 4, and model 5 was not obviously changed compared to that of model 2. As a result, model 2 was selected as the final model in the following analysis.

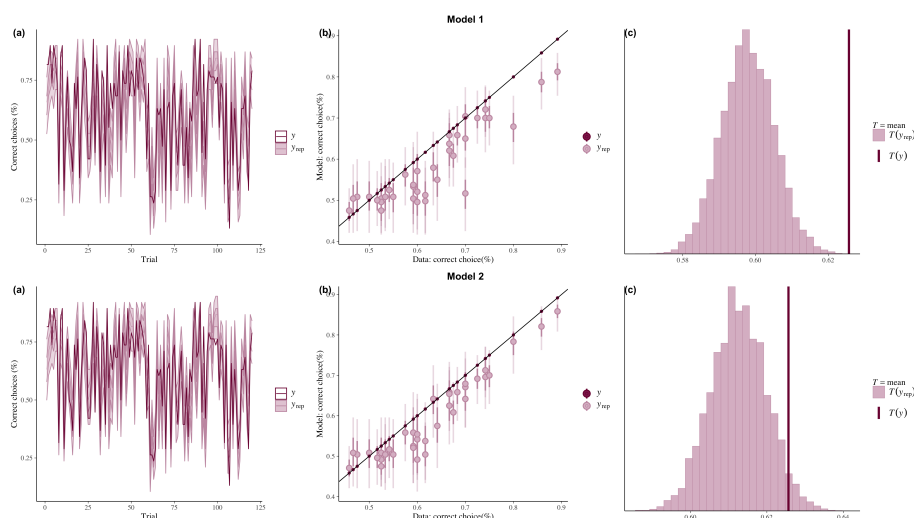


Figure 5.2: Model validation with PPC on the model-independent measure of the probability of choosing the correct choice for model 1 and model 2. (a) Trial-level model predictions plotted against actual data. Shaded area depicts the 95% HDI of the posterior distribution. (b) Subject-level model predictions compared with actual data, in relation to the identity line. Error bars depicted the 95% HDI of the posterior distribution. (c) Grand average model prediction across trials and participants.

Effects of Experimental Manipulation on Learning Rate

Having selected model 2, we fitted this model to participants' choice behavior and estimated distributions over model parameters at both individual- and population-level. Here, we used the posterior distributions over the group mean for each learning rate component to examine whether learning rate varied as a function of block type (volatile or stable), and relative outcome value (i.e. trials following good or bad outcomes). Figure 5.3 shows the posterior means, along with their 95%-HDIs for each parameter component. If the 95% HDI for a given effect does not cross zero, the effect is considered statistically credible. The effects of block type and relative outcome value upon learning rate were

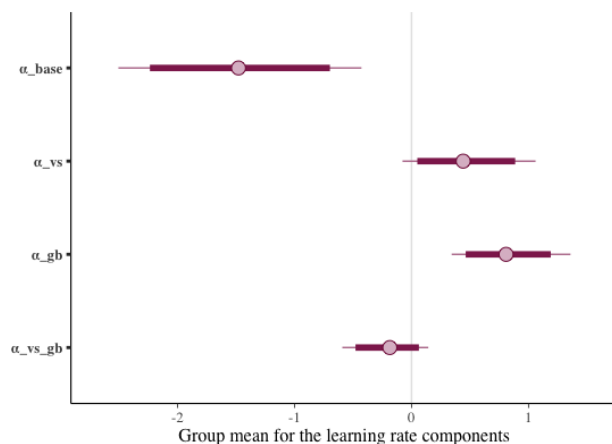


Figure 5.3: Effects of block type (volatile vs. stable), and relative outcome value (good vs. bad) on learning rate ($n = 38$). It shows the posterior means along with the 95% HDI for the group means (μ) for each learning rate component (i.e. for baseline learning rate and the change in learning rate as a function of each within-subject factor and their two-way interactions). The 95% posterior intervals excluded zero for the effect of block type upon learning rate (i.e. the difference in learning rate for the volatile versus stable task blocks α_{vs}). This was also true for the effect of relative outcome value, that is, whether learning followed a relatively good (no stealing) or relatively bad (being stolen) outcome (α_{gb}). Participants showed higher learning rates during the volatile block than in the stable block and on trials following good versus bad outcomes.

statistically credible; that is, their HDIs did not cross zero (block type α_{vs} , $\mu = 0.43$, 95%-HDI = [0.03, 0.87]; relative outcome value α_{gb} , $\mu = 0.81$, 95%-HDI = [0.47, 1.17]). Participants had higher learning rates during the volatile block than in the stable block and higher learning rates on trials following good versus bad outcomes. The two-way interactions of the two variables were not statistically credible ($\mu = -0.19$, 95%-HDI = [-0.46, 0.07]).

The Correspondence between Model-derived Parameters and Model-agnostic Measures

To broadly examine the relationships between model-derived parameters and model-agnostic measures of behavior, individual estimates of the mean values for each parameter were extracted and compared to the model-agnostic behavioral measures. These analyses confirmed expected relationships between these model-based and model-agnostic measures (see Figure 5.4). The difference of the proportion of choosing the correct choice ($\text{diff}_{correct}$) and the difference of the proportion of loss-shift ($\text{diff}_{loss-shift}$) were both strongly positively

correlated with the learning rate component ($\alpha_{v,s}$) that represents the difference in learning rate between the stable and volatile blocks ($\text{diff}_{correct}$: $r = 0.52, p < 0.001$; $\text{diff}_{loss-shift}$: $r = 0.53, p < 0.001$). This result was as expected because these two behavioral metrics reflect to which extent the participants updated their learning strategy when the contingency volatility was changed, which was similar to the meaning of $\alpha_{v,s}$. The overall proportion of choosing the smaller choice was strongly negatively correlated with the risk preference parameter γ that was assumed to be the same in the stable and the volatile block. The larger the γ was, the more weight the participants put on the outcome probability compared to the potential loss magnitude, thus, it makes sense that the fewer participants chose options with smaller losses.

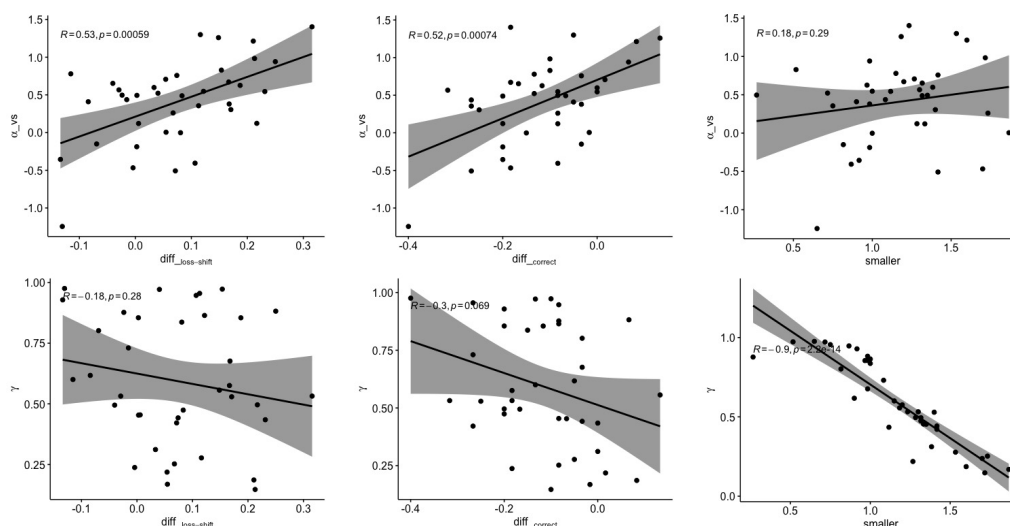


Figure 5.4: Relationship of model parameters with model-agnostic summaries of behavior.

Impacts of Non-motor Symptoms on Learning Patterns

In order to investigate if the presence of deficits in adaptive learning was affected by PD or the non-motor symptoms in PD, we explored whether the difference in learning rate between the volatile and stable blocks varied as a function of the PD severity or its non-motor symptoms. Examining learning rate difference between blocks $\alpha_{v,s}$, the 95%-HDI both for the depression and anxiety regression coefficient excluded zero (Figure 5.5a, $\beta_d = -0.48$, 95%-HDI = $[-0.88, -0.07]$; Figure 5.5b, $\beta_a = -0.46$, 95%-HDI = $[-0.90, -0.05]$). Individuals with lower levels of depressive and anxiety symptoms adjusted their learning

rates between the stable and volatile task blocks to a greater extent than individuals with higher levels of depressive and anxiety symptoms. Both the estimates for the main effect of relative outcome value (α_{gb} ; $\beta_d = -0.03$, 95%-HDI= $[-0.46, 0.40]$; $\beta_a = -0.28$, 95%-HDI= $[-0.62, 0.05]$) and the interaction between block type and relative outcome were not credibly modulated by depression or anxiety symptoms ($\alpha_{vs \times gb}$; $\beta_d = -0.15$, 95%-HDI= $[-0.47, 0.15]$; $\beta_a = 0.05$, 95%-HDI= $[-0.22, 0.30]$).

The other two non-motor symptoms, apathy and pain were not associated with the adaptation of learning rate (Figure 5.5c, $\beta_{ap} = -0.67$, 95%-HDI= $[-0.81, 0.11]$; Figure 5.5d, $\beta_p = -0.46$, 95%-HDI= $[-0.90, -0.05]$), although the apathy severity, like depression and anxiety, was associated with increased risk preference value (γ ; $\beta_d = 0.77$, 95%-HDI= $[0.01, 1.60]$; $\beta_a = 0.88$, 95%-HDI= $[0.09, 1.69]$; $\beta_{ap} = 0.99$, 95%-HDI= $[0.21, 1.81]$). It was also found that the measures that the severity of PD did not contribute to the less extent of adjustment of learning to volatility and it did not alter the risk preference parameter either (Figure 5.5e).

5.4 Discussion

We examined how patients with PD performed in terms of adaptation of learning rates to contingency volatility in a lab-in-the-field design and if their performance could be attributed to the severity of the PD or the non-motor symptoms in PD in this case study. We modelled participants' performance on a web-based aversive version of probabilistic decision-making under volatility task using the hierarchical Bayesian framework. Our results suggested that overall, patients with PD, like the healthy populations without PD reported in other studies [103], were able to rationally adapt their learning about aversive outcomes on the basis of whether action-outcome associations were volatile or stable. The effects of the variables of interest were examined within the model by incorporating other population-level parameters that captured variance attributable to each of the covariates. The severity of the depression and anxiety symptoms in patients with PD were uniquely associated with impoverished adaptation of learning to volatility among the other non-motor symptoms and the PD itself, but this association was not modulated by the adjustment

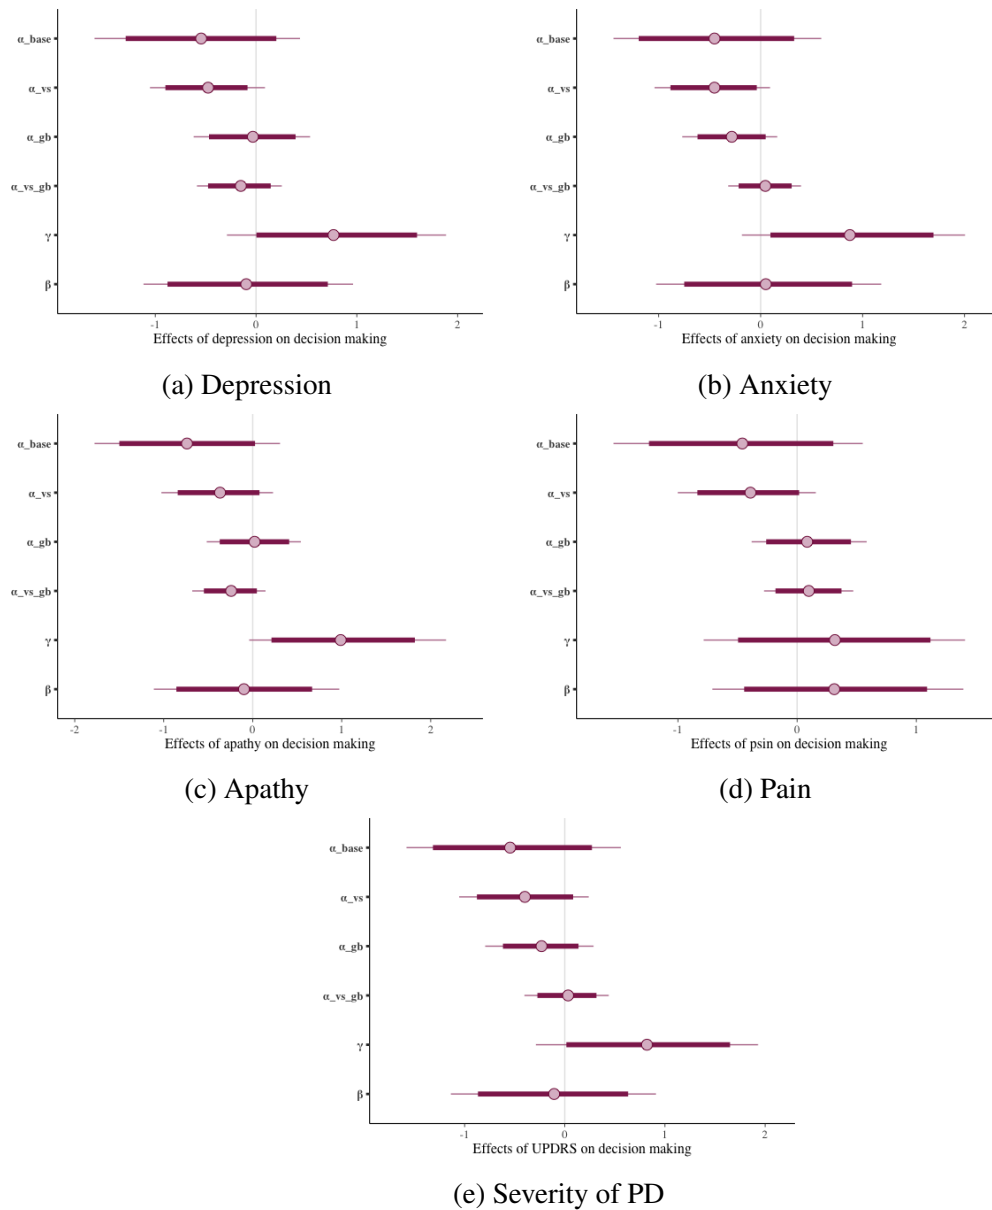


Figure 5.5: Effects of the severity of PD and the severity of non-motor symptoms of PD on the cognitive parameters. The posterior means and 95% HDI for the effect of depression (β_d), anxiety (β_a), apathy (β_{ap}), pain (β_p) symptoms, and the severity of PD (β_u) on each of the cognitive parameters were shown respectively. The depressive and anxiety symptoms uniquely and credibly modulated the extent to which learning rate varied between the stable and volatile task block (α_{vs} ; $\beta_d = -0.48$, 95%-HDI = $[-0.88, -0.07]$; $\beta_a = -0.46$, 95%-HDI = $[-0.90, -0.05]$). The depressive, anxiety and apathy symptoms were all significant correlated with the risk preference parameter (γ ; $\beta_d = 0.77$, 95%-HDI = $[0.01, 1.60]$; $\beta_a = 0.88$, 95%-HDI = $[0.09, 1.69]$; $\beta_{ap} = 0.99$, 95%-HDI = $[0.21, 1.81]$).

to relative outcome domain. This finding is in line with the existing findings observed in the general population without PD, in which decreased ability to appropriately adjust updating of outcome expectancies between stable and volatile environments represented a common vulnerability linked to both anxiety and depression [104], [198], [211]. A more thorough understanding of the relationship between learning and decision-making and the neuropsychiatric symptoms, like depression and anxiety, will provide greater insight into the nature of decision-making disturbances in patients with PD and influence new strategies to help people with this condition. Thus, our findings may facilitate early detection and add future tailored therapeutic interventions and thereby promote accurate treatment of these two neuropsychiatric disorders comorbid to PD. Since the altered computational predictors of the deficits have been linked to the activity levels in anterior cingulate [103] and broader frontal-striatal circuits [212] in neuroimaging studies, the results of this study are also useful in terms of enhancing our understanding to the pathology of these two symptoms in PD.

Although with discrepancies, it has been reported that patients with PD were markedly deficient with respect to the integrity of probabilistic reversal learning, which is comparable to the volatile block of the ALT used in this study, in patients with PD ON/OFF medication [178], [182], [188]. Most of them did not take the non-motor symptoms, such as anxiety and depression, which are potential confounding factors contributing to the learning deficits, into consideration and none of them has detected particular learning and decision-making impairments that are specifically attributed to the non-motor symptoms. Given that the decreased adaptation of learning to volatility was credibly explained by the severity of the depression and anxiety symptoms of patients with PD instead of the disease itself in this study, the discrepancies of the reversal learning deficits in PD might derive from the differences in terms of the severity level of depression and anxiety among the patients. It is necessary to exclude patients with such psychiatric disorders when exploring the influence of dopaminergic medication on feedback learning in future studies.

An extra finding of this study was that the severity of depression, anxiety, and apathy symptoms in patients with PD were associated with the risk preference parameter. The

higher the symptom scores were, the more the participant depended on the outcome probabilities of the alternatives compared to the loss magnitudes. In other words, the patients with a higher level of depression and anxiety not only were less able to adjust their learning rates to obtain accurate estimations of the loss probability of each option, but they were also more reliant on the wrong estimated probabilities when integrating the probability and magnitude information. An inability to adapt learning rate to current levels of volatility and an elevated reliance on the wrong estimated probabilities collaboratively resulted in patients with PD with higher depression and anxiety levels being less able to determine the best course of action when faced with unexpected outcomes, thus leading to poor decision-making. We did not identify similar effects of severity of PD itself or other non-motor symptoms on the adaptation of learning to contingency volatility and risk preference.

The measures of learning pattern and risk preference that are independent of formal computational models were also investigated by assessing the correspondence between the model-derived parameters and model-agnostic measures of behavior and the relationships between the model-agnostic measures and non-motor symptoms of PD. Patients selected a smaller number of correct options and increased their loss-shift rate in volatile blocks relative to stable blocks. This is consistent with the results of the computational modelling analysis that revealed a significantly elevated learning rate for the volatile versus stable task blocks. In fact, the increase of the loss-shift strategy adopted in the volatile relative to the stable block significantly corresponded to the elevated learning component that represents the difference in learning rate between the stable and volatile blocks. Nevertheless, unlike the strong effects of depressive and anxiety symptoms on the computational parameters that represent the learning strategies, relative weak effects were identified upon the superficial behavioral-level change of loss-shift rate. The proportion of choosing options associated with smaller magnitudes was constant across the stable and volatile blocks, indicating the risk preference was not affected by the volatility of the environment. The preference for options with smaller loss magnitudes was significantly correlated with the risk preference parameter, i.e. the larger the risk preference parameter, the less weight was put on loss

magnitudes, and the less the participant choose the options with smaller magnitudes. Combining the results from the computational modelling analysis, patients with a higher level of depression, anxiety and apathy level demonstrated a tendency for risk-seeking, although, unlike the risk preference parameter, the behavioral preference for smaller losses was not correlated with any of the non-motor symptoms of PD. These analyses confirmed the expected relationships between these model-based and model-agnostic measures. The diminished effects of the depressive and anxiety symptoms on the traditional behavioral summaries illustrated how the model-derived measures facilitate the detection of significant findings even with a relatively smaller sample size and noisier web-based experimental design.

5.5 Conclusion

In summary, this study identified depression- and anxiety-correlated deficits on adaptation learning to the contingency volatility in patients with PD. It is worth noting the significance of extending the lab-in-the-field methodology to this special group who may have difficulties attending to clinical interviews due to their limited mobility abilities. This case study also demonstrated the versatility of using computational models to detect potential decision-making deficiencies of various population. Like we concluded in the chronic pain case study, before using the objective decision-making measures as a daily self-checking tool for the patients, more rigorous testing and comparison in various settings is required to validate their equivalence in the lab and field.

Chapter 6

The Effects of Tinnitus in Probabilistic Learning Tasks: an EMA study

In this chapter, we implemented a proof-of-concept EMA study for the assessment of tinnitus experience and its impact on decision-making. The web-based decision-making paradigm, along with the momentary emotional states and tinnitus symptoms were repeatedly sampled at moments with high temporal resolution. Tinnitus is a disease characterized by irregular circadian rhythms and it has been reported to have negative impacts on cognition. Patients with this disease might have stochastic fluctuations in symptom and cognitive function that cannot be easily predicted. The longitudinal EMA experimental design allowed us to link the momentary computational phenotypes of decision-making to real-time psychological state and tinnitus distress. Eight tinnitus patients were recruited with the help of staffs of tinnitus clinics. The mobile app AthenaCX was applied to deliver both the baseline tinnitus and psychological assessments and the regular EMA of perceived tinnitus symptoms. The web-based ALT was triggered based on participants' responses to the EMA questions. Participants demonstrated relatively high response rates, suggesting that it is feasible to incorporate behavioral paradigms in EMA studies thus to examine if the computational phenotypes of decision-making are static traits of an individual or fluctuate with the momentary variables. The mixed effect modeling analysis revealed that, instead of being a stable feature, the primary computational phenotype of interest,

which captured participant's adjustment of learning rate to volatility, reduced when the tinnitus stressfulness was relatively higher. This result has implications to the future research in which decision-making paradigms are involved to measure human behavior. It is necessary and important to shore up equivocal between-subject effects through within-person analysis.

6.1 Introduction

6.1.1 Background

Tinnitus is an increasingly significant health concern characterized by the perception of sound in the absence of external stimuli [213]. It has been reported by the American Tinnitus Association that roughly 15% of the public, more than 50 million people in the United States, suffer from some forms of tinnitus, with 20 million people struggling with burdensome chronic tinnitus, and over 2 million characterized as extreme and debilitating cases [214]. A recent large research study conducted by Stohler et al. [215] revealed an increasing incidence rate of tinnitus between 2000 and 2016. It was reported that the number of people living with chronic tinnitus is set to increase by more than half a million over the next decade, emphasizing a potentially increasing burden on the health care system. Particularly, the challenges associated with responses to the COVID-19 pandemic increased tinnitus distress in case people perceive the situation as generally stressful with increasing grief, frustration, stress and nervousness [216]. In fact, the British Tinnitus Association has reported a rapid increase in the number of people accessing their services, with a 256% increase in the number of web chats from May to December 2020 compared with the same period in 2019.

6.1.2 Tinnitus and Cognitive Impairments

Although most tinnitus patients can cope well with the condition, managing to minimize its impact on their life, approximately 20% of individuals can be characterized as being severely debilitated by their symptoms [217], [218]. Recently it has been proposed to

differentiate between "tinnitus" to describe the auditory phantom perception and "tinnitus disorder" for the description of the auditory component plus the associated suffering [219]. Although neuroimaging evidence is emerging that tinnitus is associated with abnormal functioning of the central auditory system [218], [220], [221], epidemiological studies have revealed that the perceived sound typically associated with the condition is not the only symptom. This suggests other pathological elements may be associated with the condition. For example, the experience of tinnitus is related to a significant decline in cognitive functions, such as working memory and attention [222], [223], learning [224], and cognitive speed [225], leading to an obvious decrease in quality of life. This involvement of non-auditory impairment is reflected by the fact that tinnitus is not only related to abnormal functioning in auditory, but also in non-auditory brain areas, for example the prefrontal cortex and the anterior cingulate cortex [226], which play crucial roles in learning and decision-making [227]. Earlier studies that investigated cognitive impairments caused by tinnitus were mainly in the domains of attentional process and memory bias and the findings were largely based on patients' self-report behavioral and emotional responses to neuropsychological tests [224], [228]. Andersson et al. [229], [230] were among the first to adopt experimental techniques from cognitive psychology, i.e. the Stroop test, to measure selective attention in this context. Subsequent studies using similar methodologies further corroborated their findings that tinnitus depletes attention resources and results in compromised cognitive performance [225], [229], [231].

6.1.3 The Relationship between Tinnitus, Cognition and Psychological Disorders

Apart from impairments in cognition, it has been documented that tinnitus patients may suffer from a variety of psychological disorders such as depression and anxiety. Holgers et al. [231] found that the occurrence of depression and anxiety among a population suffering from severe tinnitus was significantly higher than that of the general population. Similar results were obtained in [232] where subjective tinnitus severity demonstrated a strong correlation with psychological distress measured by the Hospital Anxiety and

Depression Scale. As a result, the identification of depression and anxiety disorders is of significant importance in the management of tinnitus patients as these comorbidities should be specifically treated [233].

It is widely recognized that the relationships between tinnitus and psychological variables are complex. For example, it is unknown whether cognitive impairments are caused by severe tinnitus directly, or if psychological factors are also involved given that there is growing evidence for cognitive dysfunction in depression and anxiety. In other words, cognitive impairments among those with tinnitus may not simply be the result of tinnitus but the co-occurrence or mediation of high-level of anxiety and depression. Alternatively, tinnitus may lead to anxiety and emotional distress that, in turn, disrupts cognitive processes. Understanding the origin and underlying mechanisms of tinnitus and tinnitus related impairment is therefore a significant challenge for current basic research. Psychological factors, along with impairments in cognition, have been considered as covariates to predict self-reported tinnitus severity. Interestingly, although both were identified as significant predictors of tinnitus severity, the decline in cognition has not been explained so far by psychological covariates in a tinnitus population [229], [231].

6.1.4 Inclusion of Tinnitus Symptom Dynamics

Tinnitus is a subjective symptom, thus it is often difficult to measure it. The evaluation of tinnitus can be even harder given the fact that the perception of tinnitus loudness and distress is not constant in most cases, but varies over time [234]. It is not clear whether the previous between-participants findings capture trait-like features of tinnitus or state-like features associated with fluctuating tinnitus symptoms. Additionally, in the clinical management of tinnitus patients, it is critical to have useful tools for measuring the fluctuations of tinnitus symptoms and identifying predictors for severe tinnitus suffering. Conventional cross-section experimental designs are not optimized for such individual-level prediction or for the study of highly dynamic disorder like tinnitus with varying symptom trigger between individuals. This study employed longitudinal EMA design to explore how individual-level moment-to-moment changes in tinnitus symptoms affects

decision-making performance on the web-based ALT. Tinnitus symptoms in the current moment are sampled multiple times per day via self-report questionnaires and the decision-making task is triggered on several occasions based on their tinnitus severity. In contrast to retrospective self-report measures where patients are required to recall and summarize their tinnitus experience in the past one or two weeks, an EMA focuses on the current moment, minimizing the potential for recall bias and increasing ecological validity. The use of EMA in tinnitus studies has increased with the development of mobile applications and the growing availability of smartphones [235]–[238]. The mobile app, AthenaCX [239], introduced in **Chapter 3**, is used in this study to automatically send notifications to the participants at several time points during the day requesting that they complete a state questionnaire asking about their current tinnitus symptom levels.

In summary, we aim to delineate the within-subject correlation between tinnitus and cognitive impairments in this study, leveraging the flexible and powerful EMA tools. We hypothesize that tinnitus patients may suffer from cognitive degradation, thus impact their adaptation of learning to contingency volatility in the ALT, as quantified by the computational parameters extracted from the models constructed for this task. Given the dynamic essence of tinnitus symptoms, the impairment-level of adaptation of tinnitus patients may exhibit variability and is positively correlated with moment-to-moment tinnitus severity in the longitudinal study.

6.2 Methods

6.2.1 Recruitment

Participants were recruited through clinics and tinnitus supports groups and associations. We used an advertisement seeking for people whose lives are impacted by tinnitus. Eligible participants are between the age of 18 to 70 years, have subjective tinnitus of 6 months' duration or longer and have access to a smartphone device with internet capability. To incentivize the participants to be adherent with the requests for data and in addition to engage with the experimental decision-making task to the best of their ability, the payment

was a gift card and the value of the card was composed of two parts, i.e. there was a basic payment (€10) plus a variable bonus (up to €30). The bonus was determined by their response rates and performance in the decision-making task (measured by the remaining gold coins in each session) in the experiments. If the response rate to the EMA surveys was higher than 50%, the participant was eligible to obtain €10 bonus. The participants were required to complete at least 4 times of the decision-making task, each worth up to €5, which means they would get up to €20 for completing the decision-making task. For all participants, they needed to register their interest via the link provided in the recruitment poster. After receiving their registration, we forwarded the instructions for joining the experiment and asked each participant to meet the investigator online in case they had any questions. Once recruited, we gave the participant a unique ID and a link to download the study app AthenaCX from Google Play Store or Apple Store. Participants used the assigned ID to enter the study app. After opening the app, they were directed to read the Plain Language Statement and the Data Privacy Statement. Once they agreed to the terms, they were asked to consent to the study. All communication with interested participants was performed online. All assessments were completed inside of AthenaCX except for the web-based decision-making task. The collected data was stored in the AthenaCX databases, which resides on the Amazon Web Service platform in its Western Europe data center. We also followed up with the participants on the 4th day to check if they came across any problems of receiving notifications and to encourage them to be more positively engaged (see Figure 6.1 for the workflow of the recruitment process).

6.2.2 Study Design

The experiment started when the participant logged in to the app and consented to participate in the study. The complete study comprised two-phases of experiments, the cross-sectional and the longitudinal experiment, lasting up to 2 weeks depending on frequency of patient responses. The cross-sectional experiment took place right after they logged into the app and have consented to participate. In this experiment, the participants were required to finish several baseline assessments including the European School for

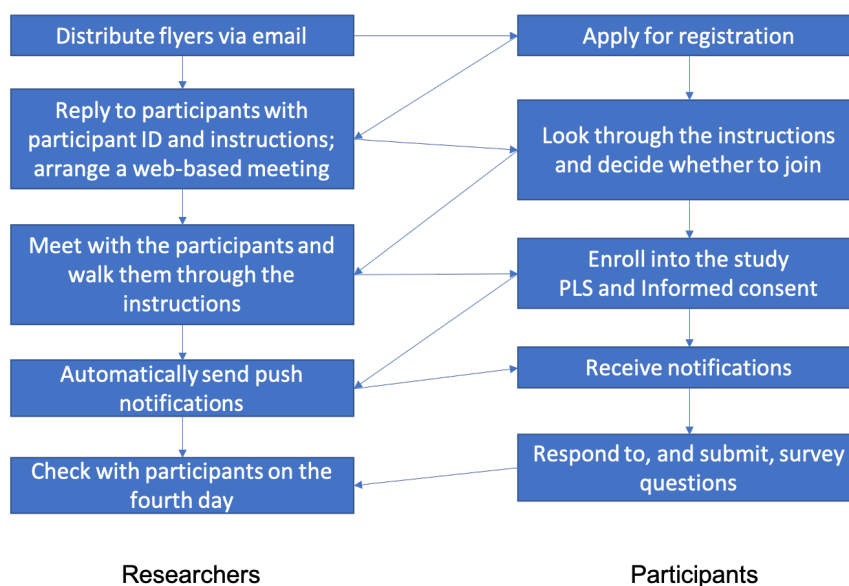


Figure 6.1: The workflow of the recruitment process.

Interdisciplinary Tinnitus Research-Screening Questionnaire (ESIT-SQ) [240] to evaluate their tinnitus-related history, State-Trait Anxiety Inventory (STAI) [241] and Major Depression Inventory (MDI) to score their anxiety and depression levels, Mini Tinnitus Questionnaire (MTQ) [242] to assess their baseline tinnitus severity, and an EMA questionnaire asking about their current tinnitus symptoms and emotional status. In summary, the participants finished three instrumental surveys and one EMA survey in the baseline assessment taking approximately 15 minutes. The last question of the EMA survey contained the link to the web-based ALT, which directed the participants to their browser. The ALT required around another 10 to 15 minutes to complete. The study was approved by the Dublin City University research ethics committee (DCUREC/2021/070).

After finishing the baseline assessments, the longitudinal-phase of the experiment was activated. AthenaCX started automatically presenting to the participants the same EMA questionnaire as they had done in the baseline experiment. It was presented from 8am to 8pm three times per day and was valid for 90 minutes. Unlike in the baseline assessment, the participants would not be presented with the ALT question until their current tinnitus symptoms reached to the predefined thresholds. Specifically, a higher and a lower threshold to the momentary tinnitus symptom was defined to trigger the ALT. However, this is unknown to the participants. They were instructed to faithfully

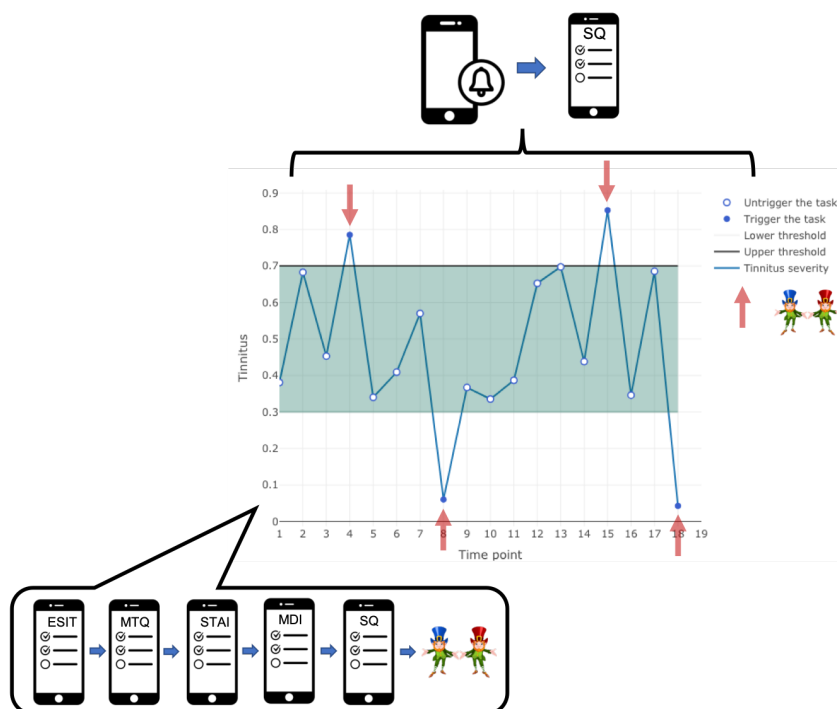


Figure 6.2: The pipeline of the experiment.

rate their momentary tinnitus symptom only. The participants were expected to complete the ALT at least four times in four different time points when their momentary tinnitus stressfulness was smaller than the lower threshold or larger than the higher threshold. If the tinnitus symptoms did not reach the threshold, the task question would not be activated and presented to the participants. Thus, we would notify the participants that they were free to withdraw from the experiment whenever the task was done four times, but they could choose to stay longer if they wished. Another termination condition was a time limit, i.e. the experiment would be expired after two weeks no matter how much data we collected from that subject. See Figure 6.2 for the pipeline of the experiments for the tinnitus participants.

6.2.3 Study App

The study app AthenaCX, which has been introduced in **Chapter 3**, is available both for Android and iOS users. All of the surveys in this study were delivered to the participants in this app.

6.2.4 The ESIT-SQ Measure

The ESIT-SQ [240] is a self-report tinnitus-relevant history questionnaire, which includes 39 multiple choice questions in total. It is structured in two parts. Part A consists of 17 questions, seven questions requiring the participant to provide details of demographics, body characteristics, education and life-style, one question related to family history, and nine questions asking for medical history and presence of hearing-related and other symptoms. Part B includes 22 questions, including eight questions about tinnitus perceptual characteristics, one general question about the impact of tinnitus, six questions about onset-related characteristics, four questions about tinnitus modulating factors and associations with co-existing conditions, one question on objective tinnitus and two healthcare-related question.

6.2.5 The MTQ Measure

The MTQ [243] is the short version of the Tinnitus Questionnaire used to examine subjective distress related to tinnitus. It consists of 12 questions reflecting the most pertinent aspects of tinnitus distress with three potential answers: true, partly true, and not true, each yielding a score from 0 to 2. The MTQ has shown good test-retest reliability, as well as high validity. The MTQ score served as the primary outcome measure for tinnitus severity in this study.

6.2.6 The STAI Measure

The STAI [241] is a commonly used measure including two, twenty-item self-report scales, for assessing trait and state anxiety, respectively. 'State anxiety' refers to the current feeling of the respondent, while 'Trait anxiety' refers to the general feeling of the respondent. Items in state scale are rated on a 4-point scale from 'not at all' to 'very much so' and items on the trait scale are rated also on a 4-point scale, but here the ordinal labels range from 'almost never' to 'almost always'. The total score of each scale ranges from 20 to 80, with higher scores indicating greater anxiety. Scores ≥ 30 indicate moderate anxiety and

scores ≥ 45 indicate severe anxiety [244]. Internal consistency coefficients for the scale have ranged from 0.86 to 0.95. Test-retest reliability of this inventory ranged from 0.65 to 0.75 over a 2-month interval [241].

6.2.7 The Major Depression Inventory

The MDI [242] is a 12-item, self-report measure for depression developed by the World Health Organization's collaborating centre in Mental health. Items contained in the MDI reflect all symptoms of depression in the DSM-IV and the ICD-10. Each of them is rated on a 6-point scale from 'at no time' to 'all the time' to assess the presence of a depressive disorder and the severity of depressive symptoms over the past two weeks. The reliability of the MDI as a measure of depression severity is 0.89 based on the results in [245]. Scores on both STAI and BDI will be used as covariates, along with tinnitus severity to examine the contribution of self-reported anxiety, depression toward the performance on the cognitive decision-making task.

6.2.8 The EMA Survey of Tinnitus Symptom and Emotional Status

The EMA survey was used to allow in-the-moment responses from participants. Most of the time, it took less than one minute to complete. It consisted of 4 questions. With the first question, we asked the participant to rate their emotional valence ranging from 0 to 10, representing very unhappy to very happy. The second question asked how much they were concentrating on the things that they were doing before the interruption. The participants could give their answers on a VAS by moving a slider between the endpoints from 'Not at all' to 'Fully concentrated'. Technically, the VAS was implemented as a slider without a pre-set initial position to avoid anchoring affects. The last two questions asked about momentary tinnitus loudness and stressfulness. The answers were also implemented with VAS. The extremes of the third VAS were 'Not audible' on one side and 'Maximal loudness' on the other side. For the last VAS the labels were 'Not stressful' on the left side and 'Maximally stressful' on the right side. We here provide the interface of the EMA questions (Figure 6.3) on the smartphone screen as an example to demonstrate the

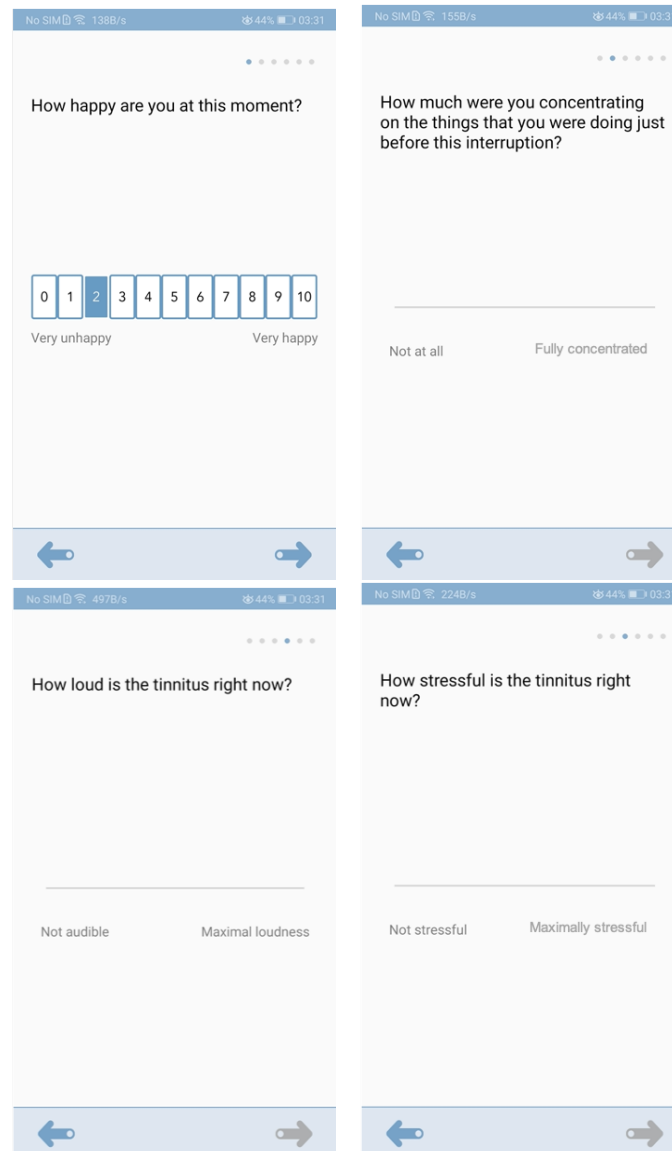


Figure 6.3: The screenshots of the interface of the EMA questions on the smartphone screen.

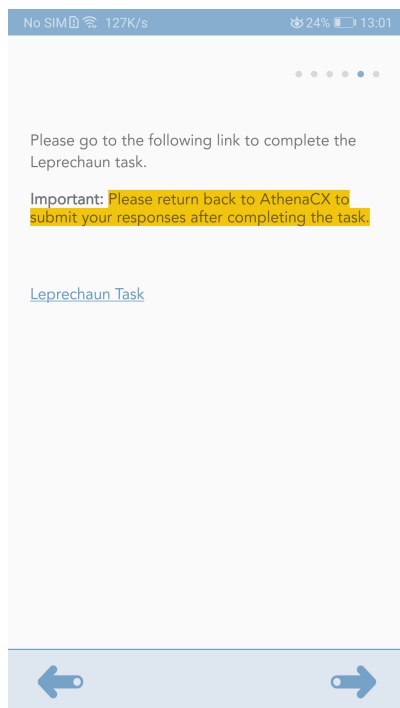


Figure 6.4: The screenshot of the task question on the smartphone screen.

implementations of the surveys on AthenaCX.

6.2.9 The Web-based Aversive Learning Task

Figure 6.4 shows the screenshot of the to-be-triggered task question based on the tinnitus stressfulness answered in last question. There was a link embedded in the question, through which the participant could be directed to the web-based ALT. The web-based ALT has been introduced in **Chapter 3**. We are not presenting more details about it here to avoid replication.

6.2.10 Compliance Analysis of the Longitudinal Part of the Experiment

To examine the feasibility of the study design, we calculated the compliance rate of the EMA survey (number of EMA prompts vs number completed) over the entire enrollment period per subject. An unconditional Linear Mixed Effects (LME) model was constructed to explore the moment-to-moment variation in the psychometric variables as assessed by the four EMA questions. We then investigated the correlations between these variables by

putting them together into the LME models.

6.2.11 Model-agnostic Behavioral Statistics

Similar to the analysis in the PD study, for each participant in each session, four basic trial-by-trial model-agnostic behavioral statistics were measured: (1) the probability of choosing the favorable option with a lower probability of stealing (correct choice), (2) the proportion of trials in which the participant did not loss gold coins (win), (3) the proportion of choosing the option with smaller magnitudes, and (4) the proportion of trials in which the participant selected a different choice after suffering from loss in previous trials (loss-shift). Each of the statistics was calculated separately for stable and volatile blocks and was submitted to repeated-measures ANOVA analysis to assess the impact of volatility of the environment on behavioral characteristics. The differences between the stable and volatile statistics were calculated to represent the adaptation of the learning and decision-making strategy on behavioral level and were utilized in the following correlation and regression analysis.

6.2.12 Group-level Computational Phenotypes across Participants and Sessions

The models used in this study were the same as those used in last chapter. Five hierarchical models were fitted and the one that was best fitted with the data was chosen for the following analysis. Data from all participants at each session was used to fit each model. Each parameter was assigned an independent group-level distribution that was shared across both participants and sessions. The primary phenotype of interest of the computational modeling analysis is if the participants adapted their learning rates in response to the change of the environment (stable to volatile block) and the outcome value (positive vs. negative results). Thus, similar to the PD study, the dependence on experimental conditions were represented through decomposing the learning rate parameter into components that captured the effects of block type (volatile, stable), relative outcome value (good, bad) and the two-way interactions of these effects.

6.2.13 Task performance Variability and Impact of EMA Variables

To examine the variation of the task performance metrics represented by the behavioral-level statistics and model-derived phenotypes, and across each metric with itself over time, pairwise correlations were run across sessions for all participants. 75% of the participants completed five times of the web-based ALT, thus the first five sessions of the task data for each participant were utilized in this analysis.

The dataset of this EMA study has a relatively complex multilevel structure. The same person answered the same self-report questionnaire and completed the same decision-making paradigm multiple sessions across the study. This means the responses were not independent. In other words, responses and task performances given by the same person would likely to be more strongly related than responses given by two different people. The non-independence characteristic of the multilevel data made it necessary to use mixed effect analysis. Thus, in order to quantify how much of the variation in task performance metrics was due to the between- vs. within-participant variability, LME models were constructed. Each behavioral statistics and session-level computational phenotypes was submitted to unconditional LME models as dependent variables without entering any predictor variables and the Intraclass Correlation Coefficient (ICC) was computed using *EMATools* package in R. To determine the degree to which variation in model-agnostic statistics and model-derived parameters could be explained by within-person vs. between-person differences in momentary tinnitus symptoms or emotional states, participant-mean levels of emotional states and tinnitus symptoms were added to the LME models as fixed effect predictors. Task repetition (number of times a subject had completed the web-based ALT up to and including the current time) was also considered as a fixed effect predictor given to examine if there was a learning effect as the repetition number of the learning task increased. LME models were estimated using *lme4* package in R. Missing data were not imputed and were censored in analyses.

6.3 Results

6.3.1 Demographic Data

Participants include 8 patients with a diagnosis of chronic tinnitus (2 females). All of the patients reported that they had tinnitus on a daily basis. Overall, the participants were moderately distressed by tinnitus, as reflected by the mean score (12.38) of the MTQ measure. The mean age of the participants was 47.25 (± 10.31) years. Detailed participants characteristics of tinnitus and psychological status are listed in Table 6.1. Tinnitus severity was positively correlated with both depression ($r = 0.68$, $P < 0.001$) and anxiety level ($r = 0.74$, $P < 0.001$), while it was negatively correlated with age ($r = -0.71$, $P < 0.001$).

Table 6.1: Basic demographic details of the participants

Number of participants (Total N)	8
Female (N)	2
Age (mean \pm sd)	47.25 \pm 10.32
MTQ (mean \pm sd)	12.37 \pm 5.40
MDI (mean \pm sd)	14.38 \pm 11.30
STAI (mean \pm sd)	42.88 \pm 14.46

6.3.2 Compliance Analysis

Participants completed a total of 202 unique EMA prompts (65.16% of all possible EMA prompts, 25.25 per person, $SD = 12.35$) over the experimental period. One of the participants completed two times of the web-based ALT, whereas all of the other participants completed more than four times (7.25 ± 5.18). Figure 6.5 (a) demonstrates tinnitus loudness and stressfulness of one sample individual varied across time points. As expected, tinnitus stressfulness varied across moments, with 44% of this variability attributed to within- person (vs. between-person) variation ($ICC = 0.56$) and tinnitus loudness also varied across moments, with 58% of this variability attributed to within-person (vs. between-person) variation ($ICC = 0.42$). Tinnitus loudness and stressfulness were strongly associated with each other (Figure 6.5(b), in which we assume the intercept varies randomly across individuals whereas the slope is the same for all individuals).

The louder the tinnitus, the more stressful the tinnitus was ($r = 0.62$, $P < 0.001$, $IC = [0.52, 0.71]$) and 52% of the variance in the tinnitus stressfulness was due to person-to-person variation, whereas 48% of the variance was due to within-person observation-to-observation variation. Tinnitus loudness also had a marginal significant negative impact on emotional state ($r = -0.07$, $P = 0.026$, $IC = [-0.14, -0.01]$), although most of the variance (66%) in the emotional state came from between-participants variation. The concentration state was not impacted by either of the momentary tinnitus-related symptoms.

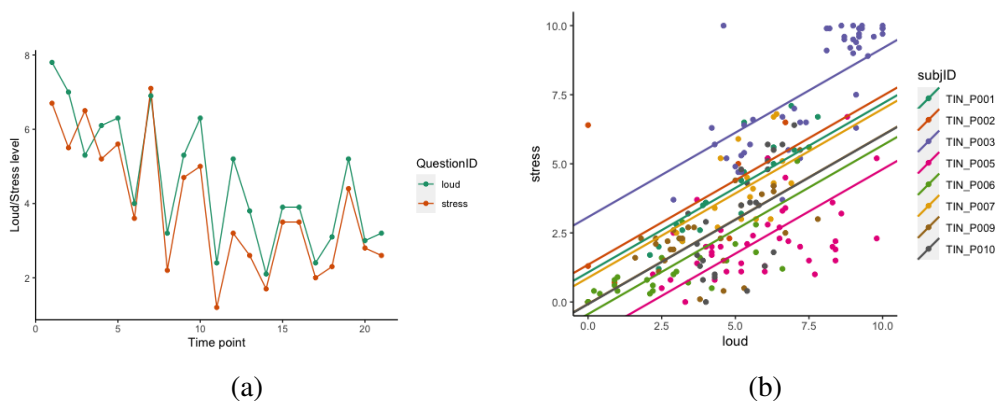


Figure 6.5: (a) Tinnitus symptom variation of one example participant. (b) Correlation between tinnitus loudness and stressfulness.

6.3.3 Model-agnostic Behavioral Statistics

Four model-agnostic behavioral statistics were summarized to obtain a visual exploration of the group-level learning patterns and decision-making characteristics (see Figure 6.6). Participants chose the option with lower probability of stealing more often in the stable block compared to in the volatile block ($F(1, 56) = 28.93$, $P < 0.001$), and win more in the stable block ($F(1, 56) = 14.27$, $P < 0.001$), regardless of the order in which the two blocks were completed (correct choice: $F(1, 56) = 0.18$, $P = 0.68$; win: $F(1, 56) = 0.05$, $P = 0.82$). Unlike the PD patients, tinnitus participants in this study did not demonstrate significant higher loss shift rate in volatile block relative to stable block ($F(1, 56) = 0.72$, $P = 0.40$), indicating that they were not more sensitive to the outcomes in the volatile than the stable environment. The preference for smaller choices was not significantly changed when the volatility of the environment was changed.

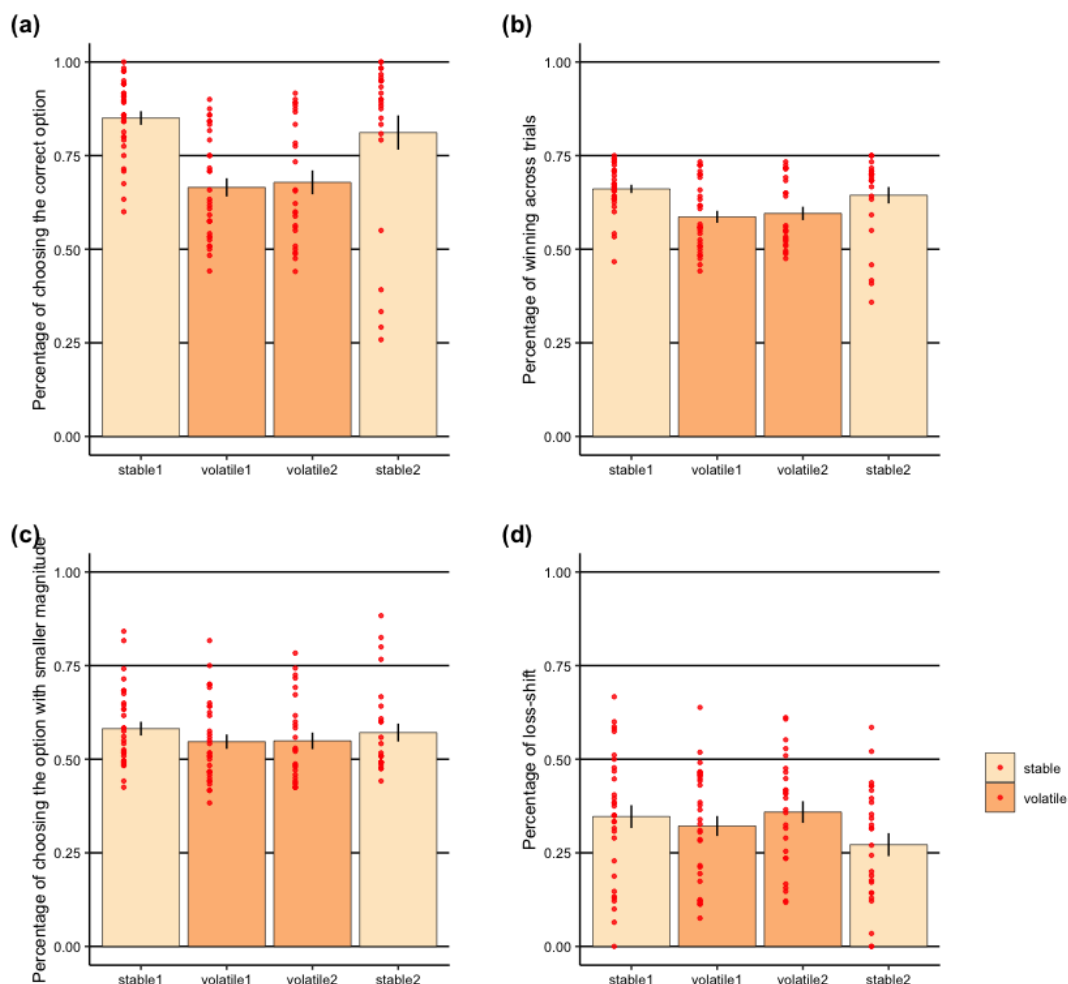


Figure 6.6: (a) The proportion of trials choosing the option with a lower probability of yielding losses, (b) the proportion of winning trials, (c) the proportion of choosing the option with smaller magnitude, and (d) the proportion of trials shifting when receiving losses for stable and volatile blocks separately for the two task schedules (schedule 1 = stable block first; schedule 2 = volatile block first). Each dot represents one session of one participant and error bars represent 1 SEM.




Figure 6.7: Effects of block type (volatile vs. stable), and relative outcome value (good vs. bad) on learning rate ($sessions = 58$). The 95% posterior intervals excluded zero for the effect of outcome value upon learning rate (α_{gb}), i.e. whether learning followed a relatively good (no stealing) or relatively bad outcome (being stolen). Participants showed higher learning rates on trials following good versus bad outcomes. However, they did not improve their learning rates in volatile block versus in stable block.

6.3.4 Group-level Computational Phenotypes across Participants and Sessions

Five models were fitted and compared in this section. The PSIS-LOO values for each model are listed in Table 6.2. Consistent with the model comparison results in the PD study, additive model was significantly better than the multiplicative model (model 1 PSIS-LOO = 10130 vs. model 2 PSIS-LOO = 9099; difference in PSIS-LOO = 1030.8, standard error of difference = 90.8). Model 3 to model 5 in which additional parameters were added did not significantly improve model fit (the PSIS-LOO decrease of these models compared to model 2 was not larger than four times of the standard error of the difference). Thus, model 2 was regarded as the best fitting model and was used in the following analysis.

Table 6.2: Leave-One-Out Information Criteria of the Five Models

Model	LOOIC
model 1	10130
model 2	9099
model 3	9026
model 4	9108
model 5	9020

Figure 6.7 shows the posterior group means (across participants and sessions), along

with their 95% HDIs for each learning rate component. The 95% HDI of the α_{gb} parameter, which represents the effect of relative outcome value upon learning rate, did not cross zero, thus was considered to be statistically credible ($\mu = 0.97$, 95%-HDI = [0.78, 1.17]). However, unlike the patients with PD, tinnitus population in this study did not adopt higher learning rates during the volatile block than in the stable block (α_{vs} , $\mu = -0.10$, 95%-HDI = [-0.46, 0.30]). The two-way interactions of the block type and outcome value were not statistically credible either.

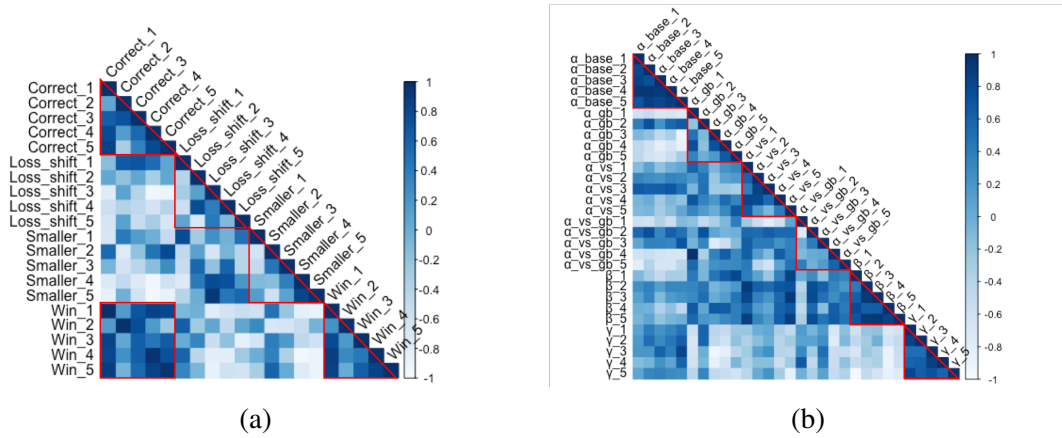


Figure 6.8: (a) Pairwise correlations between behavioral statistics of task performance across sessions and participants. The percentage of choosing the option with smaller stealing probability (*correct*) and the percentage of winning were strongly correlated *win*. These two metrics were moderately correlated with themselves across sessions. (b) Pairwise correlation between model-derived phenotypes of task performance across sessions and participants. The baseline learning rate α_{base} , the risk preference γ , and the inverse temperature parameter β were strongly correlated with themselves, whereas the subcomponents of the learning rate did not maintain constant over time.

6.3.5 Task Performance Variability and Impacts of EMA variables

Figure 6.8 shows the pairwise correlations for the (a) behavioral statistics and (b) model-derived parameters across sessions and participants, respectively. The numbers of the row and column annotations represent the number of sessions the ALT was completed. Comparing the behavioral statistics across sessions, we found the percentage of choosing the option with smaller stealing probability and the percentage of winning metrics were moderately correlated with themselves over sessions (correct: mean $r = 0.58$, $SD = 0.36$; win: mean $r = 0.62$, $SD = 0.34$) and as expected, these two metrics were moderately

correlated with each other (mean $r = 0.58$, $SD = 0.36$).

For the model-derived phenotypes, the baseline learning rate (mean $r = 0.88$, $SD = 0.09$), the risk preference (mean $r = 0.74$, $SD = 0.18$), and the inverse temperature parameter (mean $r = 0.80$, $SD = 0.17$) were strongly correlated with themselves over sessions, suggesting a degree of test-retest stability in these parameters over the sessions. The learning rate component, which represents the difference in learning rates between the stable and the volatile blocks α_{vs} , was moderately correlated with itself across sessions (mean $r = 0.62$, $SD = 0.33$). In contrast, the other two learning rate components did not keep constant at different time points (α_{gb} : mean $r = 0.42$, $SD = 0.44$; $\alpha_{vs \times gb}$ mean $r = 0.26$, $SD = 0.50$). Although the block effect learning rate component was weakly correlated with the inverse temperature parameter (mean $r = 0.51$, $SD = 0.31$), no significant correlation was identified between and with other parameters. Collectively, these data indicated that the different computational parameters assessed captured at least partly distinct aspects of participants' decision-making behavior in this probabilistic learning paradigm.

We formally quantified how much the variation in these task performance metrics was due to between-participant vs. within-participant variability using the ICC. The variation of all the model-agnostic behavioral statistics were mainly attributed to within-participant variation. The highest between-participant ICC was 26% for the percentage of winning metric. For the model-derived baseline learning rate (α_{base}) and the inverse temperature (β), more than 60% of the variation were attributed to between-participant variation, whereas the outcome value effect learning rate component (α_{gb}) and the interaction effect learning rate component ($\alpha_{vs \times gb}$) were most influenced by within-participant factors (92% and 91% respectively). 64% of the variation in the block effect learning rate component and 57% of the variation in the risk preference could be attributed to within-participant variation.

Finally, we examined the relationship between fluctuating emotional and tinnitus states and the task performance metrics as a potential source of within-participant variation. It was found that the momentary tinnitus stressfulness modulated the extent to which learning rate varied between the stable and volatile blocks ($t = -5.61$, $B = -0.18$, 95%

$CI = [-0.25, -0.12]$), with most of the variation explained by within-participant variation. Adding the individual-level psychometric assessments (such as depression, anxiety, baseline tinnitus severity, etc) as predictors to the LME models, no credible effects were identified on the task performance metrics. These results suggested that participants adjusted learning rate between the stable and volatile task blocks to a greater extent when they were under lower levels of tinnitus stress compared to higher levels of tinnitus stress. It is worth noting that none of the task performance metric was predicted by the number of times the task was completed, implying that the learning feature of the ALT was not a problem for the repeated-measure design.

6.4 Discussion

In this study, a proof-of-concept EMA study was presented for high-frequency probes of the momentary emotional and tinnitus states and performance on a decision-making paradigm requiring adaptive learning, aiming to examine the feasibility of such experimental design and explore if the computational phenotypes of decision-making are traits of an individual or fluctuate with the momentary emotional states or tinnitus-related distress.

Tinnitus patients recruited in this study demonstrated relatively high compliance in terms of using smartphone to submit information about various aspects of tinnitus bother and completing the time-consuming decision-making task when it was triggered. This observation suggests that smartphone, or the AthenaCX app used in this study, is feasible for the longitudinal momentary assessment of tinnitus perception and the enhancement of capturing the dynamic changes in probabilistic decision-making performance in the everyday life. Compliance for EMA prompts were acceptable given that most studies have not attempted to probe multiple times of probabilistic decision-making and with high temporal resolution, especially in a patient sample. Although few longitudinal studies have been conducted to investigate the temporal performance on preference-based decision-making paradigms in opioid addiction [246], [247], to the best knowledge of us, no one has examined the feasibility of dense sampling the learning-based decision-making paradigms. The main concern about using learning-based paradigms in longitudinal

experimental design is the learning effect, i.e. people might be able to remember which option yields rewards and punishments after several times of repeating of the task. In the field of computational psychiatry, decision-making based on RL from feedbacks has been most widely used for identifying disrupted process in individuals with various psychiatric disorders, however, most in cross-sectional studies. The fact that the number of task repetition did not impact on participants' task performance in this study provides experimental evidence on the feasibility of adopting repeated-measures design for learning-based decision-making paradigms to enrich our understanding to the underlying disruptions of mental disorders.

Consistent with previous findings in large-population based studies [248], which reported elevated prevalence of depression and anxiety among tinnitus patients, we revealed a strong association among tinnitus severity, depression and anxiety. The momentary tinnitus loudness and stressfulness varied considerably between patients and within individual patients including obvious fluctuation over the course of different time points [238] and the louder the tinnitus was, the more stressful it was. Although it has been documented that healthy subjects were able to assess volatility in an optimal manner and adjust decision-making accordingly [103], tinnitus patients in this study did not manage to adjust their learning rates based on the changes of the volatility of the environment across all sessions and individuals, no matter according to the behavioral statistics or the computational phenotypes. A matched healthy control group is needed in future studies to confirm if this is a specific deficit suffered by tinnitus patients. The individual-level depression and anxiety traits were not found to be correlated with impaired adaptation of learning to contingency volatility, which is as expected given the small sample size. Instead, we found that this deficit was associated with elevated momentary tinnitus stress level. In other words, tinnitus patients were less able to adapt their learning rates based on the volatility of the environments when they were experiencing higher tinnitus stressfulness. This implies that within-person changes in tinnitus symptoms may influence the ability of adaptive learning which may be consequently related to changes in broader adaptive learning behavior. The findings in the exploratory analysis of the computational parameters also backed up these

regression results because most of the variation in the block effect learning rate component were attributed to within-person variation, indicating that this phenotype might be moment sensitive. The variations in all of the behavioral-level statistics and the outcome value effect learning rate component were also mainly from within-person variation, although we did not observe additional relationships between the assessed momentary variables and the other task performance metrics. These findings highlighted the importance of complementing between-subjects assessments with more granular within-subject probes to provide a more nuanced view of how tinnitus symptoms and their impacts on decision-making interact on different timescales and levels of abstraction.

6.5 Conclusion

In this study, the momentary stressfulness of tinnitus was found to have a negative impact on participant' adaptation of learning to contingency volatility, highlighting the necessity of detecting the within-person variations of computational phenotypes and their associations with momentary psychological or physical states. Longitudinal EMA designs, such as the one used in this study, could present significant measurement challenges. These challenges require not only consistency in the effects observed at the group level, but also reliability at the individual level in order to accurately measure meaningful variations in the mechanisms. However, the extension of computational phenotyping of human decision-making to individual level prediction and longitudinal within-person change is a critical pathway forward to realizing the goals of computation to clinic translation. Thus, the successful implementation of this methodology in this study, although only within a small scale, has implications to many other behavioral and cognitive studies, providing empirical evidence for the feasibility of dense sampling behavioral paradigms based on momentary clinical symptoms that are featured with dynamic fluctuations using smartphones. The advancement of measuring momentary clinical symptoms has the potential to reshape clinical decision-making, e.g. studies examining treatments or interventions with high temporal resolution may derive novel treatment targets.

Chapter 7

A Comparison of Distributed Machine Learning Methods for the Support of Large-scale Computational Phenotyping of Decision-making

This chapter presents the feasibility of applying distributed training techniques for conducting large-scale computational phenotyping of human-decision making, which is an essential step in clinical trials. In practice, large-scale phenotyping involving human subjects is hindered by limited dataset availability for model training and validation, due to the strict legal and ethical regulations for large-scale data collection and data exchange between organizations in order to protect personal privacy. Distributed machine learning techniques, especially Federated Learning (FL), which was designed for the preservation of data privacy, is a promising method for addressing the issues of data availability. To verify the reliability and feasibility of applying FL to train DL models used in the characterization of decision-making, we conducted experiments on a real-world, "many-labs" data pool including experiment data-sets from ten independent studies. FL was compared with other distributed learning frameworks and the Centralized Learning (CL) framework in which the researchers need to have full access to the data. The results demonstrated

that the model trained using FL approach achieved performance comparable to that of the CL strategy. This suggests FL has value in scaling data science methods to data collected in computational modeling contexts when large-scale data collection or data sharing is not convenient, practical or permissible.

7.1 Introduction

The successful implementations of the case studies introduced in the preceding chapters (**Chapter 4** to **Chapter 6**) suggest that it is feasible to measure subjects' decision-making performance outside of the lab and even to sample them multiple times with a longitudinal experimental design. The strategy of allowing subjects to provide study data using their own internet enabled devices (e.g. smartphones) in these studies makes it much easier to conduct inexpensive large-scale computational phenotyping of human decision-making study. Ideally, the self-report measurements and decision-making tests can be rolled out at scale to hundreds and thousands participants simultaneously and through time. Apart from collecting positive data, smartphones support the integration of passive data generated automatically without requiring active engagement or submission of data by the participant. These features allow researchers to obtain substantially richer, multivariate data and larger samples in a short time with minimum cost.

However, large scale data collection from participants' smartphones is subject to strict laws and regulations to protect personal privacy. Both the United States Health Insurance Portability and Accountability Act (HIPAA) and the European General Data Protection Regulation (GDPR) [249] have enforced stricter regulations of the storage and exchange of personally identifiable data and data concerning health. For example, apart from passing through the general ethical application for involving human participants, additional and more complicated procedures need to be addressed if a study involves large-scale processing of personal data, not to mention the significant costs incurred with movement and storage of large amounts of data. When passive data is involved, even more ethical issues need to be taken into consideration. Passive data is often collected without the participants' active engagement or knowledge, causing greater risks to privacy and

informed consent than active data collection. An alternative way to increase data size and diversity is through collaborative learning or meta-analyses, in which many laboratories cooperate and share data to train a global model together. Nevertheless, data sharing is not easy either due to privacy, ethical, and data regulation barriers, especially when involving larger number of laboratories from different legal jurisdictions [24]. As a result, large-scale computational phenotyping in human subjects is still limited in the literature. Combining the passive data to infer aspects of momentary decision-making behavior is in absence. The ethical and legal issues of large personal data collection and sharing are also one of the main challenges we faced for scaling up the three case studies we had done.

The field of secure and privacy-preserving AI offered techniques to help bridge the gap between personal data protection and data collection and utilization for research involving large-scale phenotyping of human behavior. FL is a distributed machine learning paradigm designed to directly address the problem of data governance and privacy. It was first introduced by Google AI [250] to allow mobile devices collaboratively learn machine learning models without sharing data from the devices. It was applied very successfully in training Google’s autocomplete keyboard application [251]. FL is a data parallel training process, in which multiple clients train a model simultaneously (each on their own data in parallel) and then send their model parameters to a central sever where these are aggregated into a global model. The central sever then sends the global model to all collaborators for further training. Each iteration of this process, i.e. parallel training, parameter update, and distribution of global parameters, is referred to as a federated round. Different from FL, the Incremental Learning (IL) [252] is another distributed learning paradigm. With this approach each client trains the model and passes the learned model parameters to next client for training on its data, until all clients have trained the model once. Another approach is an extension of IL, called Cyclic Incremental Learning (CIL). This involves the repetition of the incremental learning process multiple times, i.e. fixing the number of training iterations at each laboratory and cycling repeatedly through the clients. Chang et al. [253] explored these two distributed training approaches using medical imaging data and found the cyclic approach performed comparably to centralized training approach,

suggesting that sharing data may not always be necessary to build deep learning models for patient imaging data. Micah et al. [252], [254] compared FL and IL approaches on imaging datasets and unsurprisingly, the basic IL performed the poorest compared to FL and CIL. While CIL may seem a simpler alternative, given that implementation of FL depends on a set of key challenges [255], it required additional validation steps at the end of each cycle, which are basically as complex as the synchronization logic of FL, to achieve comparable results to FL paradigm. Critically, CIL was less stable than FL, resulting in an inferior alternative.

In previous chapters, we used computational models that specify a set of assumptions about the underlying cognitive processes of subjects to represent the phenotypes of human decision-making in a continuous parameter space. Such models generally require manual engineering with an iterative process to examine the consistency between the hypothesis and the empirical data [92]. The approach may fail if subjects adopt a completely different strategy during decision-making. DL models, especially RNN models, present an alternative approach to the characterization of decision-making behavior. They can automatically capture behavioral trends exhibited by subjects without strong assumptions about the mathematical structure of the underlying process through taking advantage of their data-driven design and higher capacity for representing complex computational processes [92], [94], [95]. Another key benefit of DL model is that they are more capable and flexible for training and testing with big data, making it a more promising and valuable tool for large-scale phenotyping of human decision-making.

In this study, we examine and compare the reliability and feasibility of applying distributed learning strategies to DL models for modeling human decision-making processes. This is demonstrated using a real-world “many-labs” dataset comprising data from 10 different laboratories [256], where we artificially introduce a restriction on sharing of data to simulate a real world example. The laboratories in this dataset can also be viewed as ‘devices’ or ‘clients’, in the context of FL. In fact, FL has emerged as a promising strategy in scenarios where, for example, hospitals operate under strict privacy practices and may face legal, administrative, or ethical constraints that require data to remain local. Subjects

from all laboratories were healthy participants and completed a computerized version of the IGT, one of the most widely used tasks measuring human decision-making under uncertainty in an experimental context. We will first train single laboratory models for each laboratory in the data pool and evaluate each of these models against held-out testing datasets from each laboratory defined prior to model training. This experiment is used to demonstrate the need for numerous and diverse data for training a robust DL model. Secondly, we will illustrate the benefits of training collaborative models involving data from all laboratories. The most straightforward way of training a collaborative model is to use the Centralized Learning (CL) strategy, that is, pooling data from different laboratories together and feeding them all to the collaborative model. However, as stated above, data sharing or large-scale data collection is not always applicable because of data protection and regulation purposes, which makes this centralized paradigm difficult to achieve in the real world. Thus, we will compare and evaluate the performances of different distributed learning strategies in which the collaborative model can be trained even without having access to the laboratory data. One centralized model and several distributed models were trained for this purpose. On-policy and off-policy simulations were then conducted to compare the performance of the three distributed learning strategies. Apart from assessing the relative performance of distributed learning compared to the CL paradigm, the simulation analyses in this study present new insights provided by DL models in terms of revealing the learning and decision-making strategies of human subjects.

7.2 Methods

7.2.1 The IGT dataset

The data pool ($N = 617$) used in this study derives from 10 studies assessing performance of healthy participants (i.e. without any known neurological impairments) on the IGT [256]. It involved a broad range of healthy populations with ages ranging from 10 to 88 with various education backgrounds and social status. Participants completed a computerized version of the IGT consisting of 95, 100 or 150 trials. All included studies used the variants

of IGT payoff scheme or the payoff scheme introduced by Bechara and Damasio [159]. Since the RNN model used for predicting subjects' actions required us to have inputs with the same size, the trial-by-trial decision sequences of all subjects were truncated according to the length of the smallest sequence, i.e. 95 trials. In other words, for those subjects who completed 100 or 150 trials of the IGT, only the first 95 trials were used for training or testing the model. Each subject was an abstract data point and each data point was a sequence of decision-making choices that contains 95 trials. The number of subjects in each laboratory, which we will refer to lab 1-10, was 15, 162, 19, 40, 70, 25, 153, 35, 57, and 41 respectively.

This dataset was reused here because, 1) participants in all studies were healthy subjects and they all completed the same decision-making task - IGT, which made it feasible to train a collaborative centralized model for all laboratories; 2) the subjects from ten studies were coming from different backgrounds with various ages and female proportions, thus, conducting multi-laboratory analysis was a natural operation to increase data size and diversity and improve the performance of the model; 3) it created a perfect real-world scenario where the local data of each lab is potentially non-IID and biased, thus it is suitable for testing out the reliability and effectiveness of distributed learning paradigms, that are sensitive to the data distribution [257]. Each laboratory can be seen as a 'device' or 'client' in the context of FL and the data assignments for the clients in FL could be matched with the real-world data distributions, such that all subjects from the same study were assigned to the same client; 4) last but not least, we can more thoroughly investigate how healthy subjects learned and behaved on the IGT using DL models, which as stated previously, are promising in modeling human decision-making because of their higher capability of learning from data in a less model-constrained manner.

7.2.2 RNN Model for the IGT

In order to model the data produced during the IGT, we used the state-of-the-art neural network approaches [92], i.e. RNNs. The RNN model is composed of a GRU layer and an output softmax layer with four nodes corresponding to the four choices for the different

decks in the IGT. The inputs to the GRU layer are the previous choices made by the subject along with the rewards and losses they received after making that choice. The softmax output are the probabilities of choosing each deck option and sum to 1. The architecture of the model is shown in Figure 7.1. It is known that the basic RNNs are inefficient in solving problems that require the learning of long-term temporal dependencies because of the gradient vanishing problem [258]. If a sequence is long enough, which is the case of our choice sequences that include at least 95 trials, they will fail to convey information (the deck preferences) from earlier trials to later ones. LSTM [93] and GRU [259] networks were created as the solution to this short-term memory issue. In most cases, the two approaches yield comparable performance. It is often the case that the tuning of hyperparameters may be more important than choosing the networks. The GRU was chosen here in this paper since it has fewer parameters and can be trained faster. We believe this to be a suitable choice given that we have a relatively small dataset with moderately long sequences.

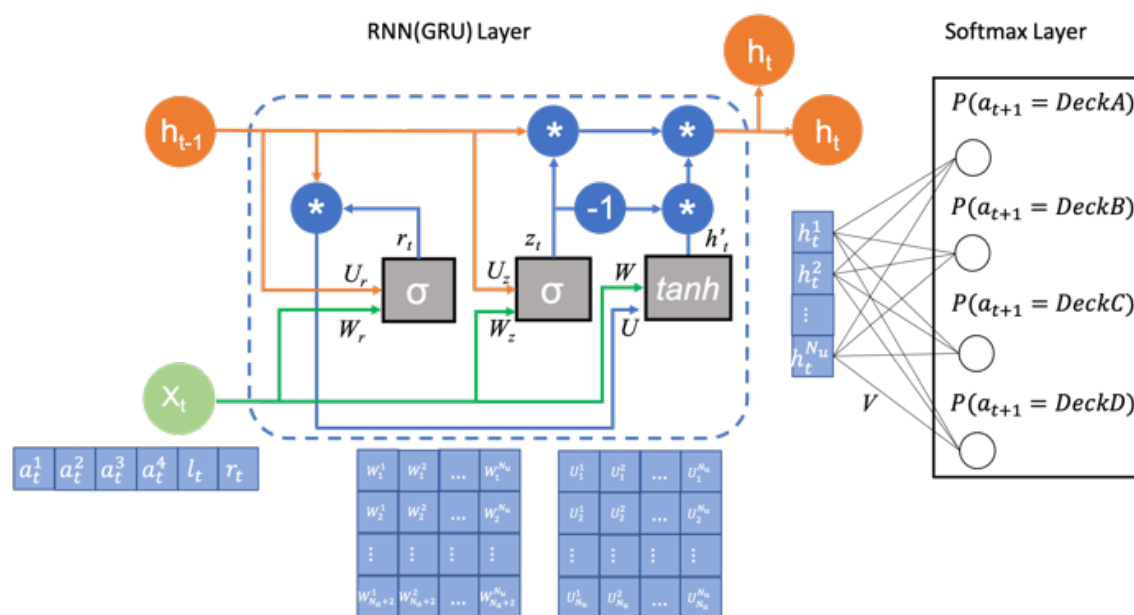


Figure 7.1: Architecture of the RNN model. The model has a GRU layer (the left part), which receives the previous action (the action is one-hot coded), loss and reward x_t and previous hidden states h_{t-1} and a softmax layer (the right part), which outputs the probability of selecting each deck on the next trial. The GRU layer is composed of two gates, i.e. an update gate and a reset gate, containing a set of weight parameters $\{W_r, U_r, b_r, W_z, U_z, b_z, W, U, b_{hW}, b_{hU}\}$ (b_r, b_z, b_{hW} and b_{hU} are not shown for simplicity). The output of the cells of the GRU layer h_t are connected to the softmax layer using black lines. The weight parameters in the softmax layer are represented as V .

The GRU layer is composed of a set of hidden units (N_u). Each unit is associated with a unit output denoted by h_t^k for cell k at time t . Let's define $h_t = [h_t^1, h_t^2, \dots, h_t^{N_u}]^T$ ($h_t \in \mathbb{R}^{N_u}$) as a vector containing all the cell outputs at time t . They are initialized with the value of zero and are updated after receiving each time step input. $x_t \in \mathbb{R}^{N_c+2}$ (N_c is the number of choices, which is 4 for the IGT) is a vector containing inputs to the network at time t , i.e. the action c_t coded using one-hot representation and the reward r_t and loss l_t received after taking action. There are two gates in the GRU network, i.e. a reset gate and update gate. The update gate helps the model to determine how much of past information needs to be passed along to the future. The output of this gate z_t is a linear combination of the input vector of the current time step x_t and the previous cell output going through the sigmoid function.

$$z_t = \sigma(W_z x_t + U_z h_{t-1} + b_z) \quad (7.1)$$

The reset gate is used to decide how much past information is forgotten. The formula is the same as that used for the update gate. The difference arises from the weights and the gate's usage, which we will examine later in this section.

$$r_t = \sigma(W_r x_t + U_r h_{t-1} + b_r) \quad (7.2)$$

The output of the reset gate is then passed through the current memory content, in which the reset gate is used to store the relevant information from the past.

$$h'_t = \tanh(W x_t + b_{hW} + r_t \odot (U h_{t-1} + b_{hU})) \quad (7.3)$$

In the last step, the network needs to calculate the h_t vector which holds information for the current unit and passes it down to the network. In order to do this, the update gate is needed. It determines what to collect from the current memory content h'_t and what to take from previous steps h_{t-1} . The parameters of the GRU layer include $W_z, W_r, W \in \mathbb{R}^{N_u \times (N_c+2)}$, $U_z, U_r, U \in \mathbb{R}^{N_u \times N_u}$, and $b_z, b_r, b_{hW}, b_{hU} \in \mathbb{R}^{N_u}$.

$$h_t = z_t \odot h_{t-1} + (1 - z_t) \odot h'_t \quad (7.4)$$

The softmax layer takes outputs from the GRU layer (h_t) as its inputs and outputs the probability of choosing each action. The parameter of the softmax layer is $V \in \mathbb{R}^{N_u \times N_a}$. The parameters of the RNN model will be $\Theta = \{V, W_z, W_r, W, U_z, U_r, U, b_z, b_r, b_{hW}, b_{hU}\}$. The RNN model was trained using the maximum likelihood estimation and Categorical Cross-Entropy Loss summed over all subjects on all trials:

$$Loss = - \sum_{s=1}^S \sum_{t=1}^T y_{s_t} \log(\hat{y}_{s_t}) \quad (7.5)$$

The model was implemented in TensorFlow [260] and optimization was based on the use of the Adam optimizer [261].

7.2.3 The Architectures of the Four Training Strategies

Figure 7.2 demonstrates the architectures of the four training strategies. The first CL strategy is the most straightforward approach. Here, data is shared among laboratories and uploaded to the center server directly. Distributed learning, i.e. the following three strategies in Figure 7.2, makes it possible to train the model through sharing the model parameters rather than the more sensitive raw data. In the Incremental Learning (IL) strategy, each laboratory trains the model once with their own data and passes the model parameters to the next client to continue the training. The training process is finished when all laboratories train the model once. In the Cyclic Incremental Learning (CIL) strategy, this process can be repeated several times. In the FL strategy, all the laboratories train the model with their data simultaneously and share their model parameters to a central server to build a global model. All laboratories will receive the parameters of the global model and update their parameters to match the global model with certain probabilities. It is expected that the CL-based model will behave the best in terms of prediction accuracy and capturing characteristics of human decision-making strategies, while the other three distributed learning paradigms will sacrifice competitiveness to different extents in exchange for

preservation of data privacy. Thus, the CL-based model will serve as the baseline model against which comparisons are made with the distributed models. All of the learning paradigms are implemented in Python 3.7.

7.3 Results

7.3.1 Experimental Settings

Each laboratory dataset was divided into training and testing datasets with 80% of the data points for the former and the remaining 20% for the latter. Subjects were not mixed across training and testing sets. Two experiments were conducted with these datasets.

- The first experiment was designed to establish the need for bigger and more diverse data to train a reliable and robust RNN model that could predict human actions with reasonable accuracy. In this experiment, we trained a single RNN model on the training dataset for each laboratory and evaluated each of these models against testing datasets from each laboratory. In total, we trained 10 independent RNN models, one for each separate laboratory, and evaluated their generalization performance both on the testing dataset of their own laboratory and testing dataset from the other laboratories.
- In the second experiment, we sought to demonstrate the improvements of collaborative models (fitted with data from all laboratories) compared to single models and also to make a comparison between the centralized learning model and the distributed learning models. The collaborative models were trained directly or indirectly with bigger and more diverse datasets depending on the collaboration methods imposed between the laboratories, i.e. CL-based, IL-based, CIL-based), and FL-based paradigms. For the IL-based and CIL-based paradigms, three different orders of training the model were performed based on the sample size of the laboratories, i.e. starting from the laboratory with the smallest sample, the largest sample or in random order, yielding three models for these two training strategies. As a result, 8

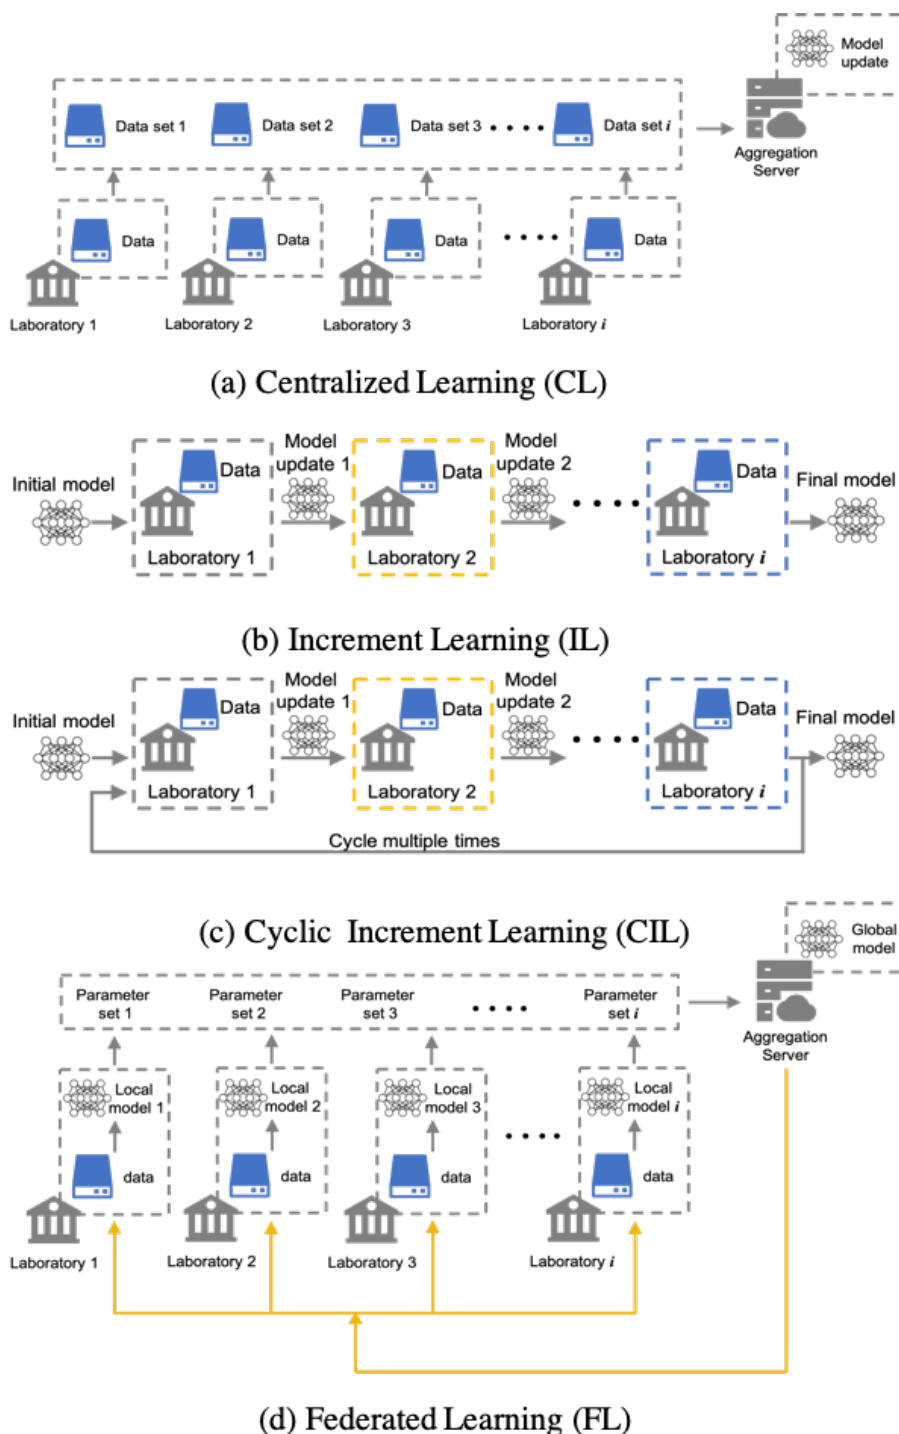


Figure 7.2: Architectures of the four training paradigms. (a) The most straightforward CL approach is based on centralized data sharing, i.e. requiring that every laboratory share data to a central node and aggregating to train a model; (b) IL approach, in which each laboratory trains the model and passes the learned model parameters to the next laboratory for training on its data, until all have trained the model once. (c) CIL, which involves the repetition of the incremental learning process multiple times. (d) FL, in which all laboratories train the same model at the same time and update their model parameters to a central sever where it is used to create a global model.

collaborative models were trained in this experiment.

The hyperparameters for training the centralized model were set as follows: batch size 32; maximum epochs 250; learning rate 0.02; number of hidden cells 10. To create an equal playing ground for the centralized learning and the distributed learning paradigms, we need to guarantee that all models see exactly the same number of samples in the whole training process. Accordingly, the number of communication and the epochs per round for the FL paradigm were set as 5 and 50 respectively. The maximum epochs for each laboratory was set as 250 for IL-based model. For the CL-based model, the frequency of weight transfer was every 50 epochs, thus the number of cycles was 5. The learning rate and the number of hidden cells were set as the same values for all models. The batch size for the distributed learning paradigms was the size of each laboratory sample. For the single models, two hyperparameters, the number of cells and epochs, were tuned via 10-fold cross validation. Please see Table 7.1, in which the parameter settings for each model are listed.

Table 7.1: Parameter settings for the centralized and distributed learning models

Model	Batch size	Number of epochs	Learning rate	Number of hidden cells	Number of cycles	communication rounds
CL-based	50	250	0.02	10	N/A	N/A
IL-based	Size of each lab training sets	250	0.02	10	N/A	N/A
CIL-based	Size of each lab training sets	50	0.02	10	5	N/A
FL-based	Size of each lab training sets	50	0.02	10	N/A	5

7.3.2 The Need for More Numerous and Diverse Data

In this experiment, ten single models were trained with the tuned hyperparameters using the training datasets from each laboratory respectively. These models were then evaluated using the testing datasets available from each laboratory. Figure 7.3 shows the testing results of the ten single models both against their own testing datasets and testing datasets




Figure 7.3: Accuracy of single lab models tested against testing datasets from each of the laboratories. The vertical axis represents models trained on a single laboratory dataset, and the horizontal axis represents the testing dataset of each independent laboratory. AVG is the average accuracy of each laboratory model performance over all laboratories. Overall, the prediction accuracy of each single model on the testing sets, no matter from their own laboratory or other laboratories, are not satisfying. Some of them are lower than the chance probability 25%.




Figure 7.4: The training and validation loss of the ten single models for each laboratory. The training loss decreases over time, whereas the validation loss starts increasing shortly after the first 10 or 20 epochs.

of other laboratories. The chance probability of predicting the next action taken by the subjects correctly was 25% since there were four options in the IGT. However, the prediction accuracy of several single models on some of the testing datasets was lower than the chance probability, which was obviously not satisfactory performance. In order to diagnose why the single models fail, we conducted a 10-fold cross validation with the training sets of each laboratory for each single model. The training and the validation loss are plotted in Figure 7.4. It can be observed that although the training loss of the single models was continuously decreasing across the training process, the validation loss started increasing shortly after around the first 10 or 20 epochs, which means the models started to overfit. It can be seen that most single models performed the best on the testing dataset of lab 9, except for the three models that trained on the laboratories with smaller sample sizes, i.e. lab 1 (12), lab 3 (15), and lab 6 (20). The single models trained on these laboratories did not perform well on their own testing sets nor did they generalize well to any other testing sets from other laboratories. The best average model performances were shown on lab 2 and lab 7, which had the largest sizes of training samples, 129 and 122, respectively. These results indicated that more numerous and diverse data were needed for a laboratory site in order to train reliable and generalizable models, at least in the contest of training RNN models.

The accuracies of single models based on testing data collected from the same lab as that which provided the training data were commonly lower than the accuracies calculated using testing data from different laboratories (e.g. the model developed using data from Lab 2 demonstrated better testing performance for the testing data associated with Lab 9 than that associated with Lab 2). This arose through chance. It may be surprising because we might assume participants from the same lab demonstrated similar behavior given that they were under similar experimental settings and thus the single models should perform better on their own datasets. However, the confounding factors associated with experimental biases should not be a significant factor and did not appear to be. It is worth considering that the testing sets (no matter where they were sourced) represented sets of distinct experimental subjects whose data had never been used in training models. In

other words, we expect all testing subjects to be independent performers of the task and model prediction on their data could contribute to an understanding of the generalization error. Testing sets from different labs contributed to our estimate of the generalization error creating a distribution. Indeed, if we observed that the testing performance was highest for the lab in which the data was collected, this might suggest that local idiosyncrasies of experimental procedures were being modelled via the data collected. This would be an undesirable phenomenon.

7.3.3 The Benefits of Training Collaborative Models

In this section, we aimed to demonstrate 1) the benefits of laboratory collaboration and 2) the feasibility of conducting distributed learning strategies when data sharing barriers were present. We compared the prediction accuracy of the collaborative models with the single models for the first purpose. The results are illustrated in Figure 7.5. As a result, apart from random and descending-ordered IL and CL models, the prediction accuracies of the collaborative models were significantly improved when testing on each laboratory's data. Notably, the average performance over all laboratories of the CL-based model were improved to 55%, 17% points higher than the lowest accuracy achieved by the single model for lab 2, and the corresponding increase for FL-based model, CIL model with ascending order and IL model with ascending order was 15%, 15% and 14% points, respectively. Since random and descending ordered IL and CIL models did not outperform the single models, only the ascending ordered models were selected to conduct further evaluation. We will refer to these as the IL-based and the CIL-based model in the following analysis for simplicity of reference. The ascending ordered IL-based and CIL-based models and the other two collaborative models were evaluated from multiple aspects for the second purpose.

7.3.4 On-policy Simulation

After training the four collaborative models, we fixed the weight parameters of the models and simulated them on-policy in the task. We fed the model with the same payoff schemes




Figure 7.5: Accuracy of four collaborative models tested against testing datasets from each of the laboratories. The vertical axis represents the collaborative models trained on the training datasets from all laboratories and the horizontal axis represents the testing dataset of each independent laboratory. AVG is the average accuracy of each collaborative model performance over all laboratories. Apart from the random and descending-ordered IL and CL model, the prediction accuracies of the collaborative models were significantly improved compared to that of the single models.

as for the subjects and the models selected the decks autonomously based on what they had learned from the human behavior data. Since 491 experimental subjects from across all the studies were used for training the model, we simulated 491 fake agents for each collaborative model.

The Average Probability of Selecting each Deck over Subjects

The SUBJ column in Figure 7.6 shows the average probability of choosing each deck over all experimental subjects. All subjects selected Deck A the least compared to other Decks, which was not surprising because Deck A was the bad deck that yields negative long-term payoff and with more frequent losses. Consistent with subjects' choices, the probability of choosing Deck A for both centralized and distributed agents was the lowest. Neither CL-based agents ($\eta = 0.005$, $SE = 0.004$, $p = 0.26$) nor the FL-based and the IL-based agents (FL-based: $\eta = 0.001$, $SE = 0.004$, $p = 0.91$, IL-based: $\eta = 0.001$, $SE = 0.004$, $p = 0.83$) were significantly different from subjects in terms of the probability of selecting Deck A, although the CIL-based agents selected slightly significant more




Figure 7.6: Probability of selecting each deck (average over subjects). SUBJ refers to the data of the experimental subjects; CL-based, FL-based, CIL-based, and IL-based are the on-policy simulations of the four collaborative models trained on the IGT with the same payoff scheme and same number of trials the subjects completed, respectively. CL-based, FL-based and IL-based agents behaved more similarly to human subjects, selecting more good decks than bad decks.

Deck A than subjects ($\eta = 0.01$, $SE = 0.005$, $p = 0.03$). FL-based agents had significant aversion to Deck D ($\eta = -0.02$, $SE = 0.010$, $p = 0.03$), while CIL-based and IL-based agents had significant preference to Deck B (CIL-based: $\eta = 0.03$, $SE = 0.009$, $p < 0.001$, IL-based: $\eta = 0.02$, $SE = 0.008$, $p = 0.010$) and significant aversion to Deck C (CIL-based: $\eta = -0.04$, $SE = 0.009$, $p < 0.001$, IL-based: $\eta = -0.03$, $SE = 0.009$, $p < 0.001$) compared to the subjects. As expected, no significant difference was observed between CL-based agents and subjects on all of these decks. This can also be seen from the plot of the overall probabilities of selecting Deck B and Deck C where IL-based and CL-based agents were relatively higher and lower than that of the subjects, respectively and the probabilities of selecting Deck D of FL-based agents were relatively lower.

It is worth noting that subjects preferred decks with infrequent losses, i.e. Deck B and Deck D, compared to decks with frequent losses, Deck A and Deck C, although they ended up selecting significantly more good decks (Deck C and Deck D) than bad decks (Deck A and Deck B) ($\eta = 0.078$, $SE = 0.010$, $p < 0.001$). CL-based, FL-based and IL-based agents managed to learn this strategy from the subjects' actions, selecting more good decks

than bad decks (CL-based: $\eta = 0.059$, $SE = 0.010$, $p < 0.001$, FL-based: $\eta = 0.059$, $SE = 0.011$, $p < 0.001$, IL-based: $\eta = 0.035$, $SE = 0.008$, $p < 0.001$), whereas CIL-based agents did not learn this characteristic of the behavioral pattern, the probability of selecting good decks was not significantly higher than that for bad decks ($\eta = -0.006$, $SE = 0.010$, $p = 0.51$).

The Fluctuation of the Probability of Selecting each Deck over Trials

The analysis in the previous section suggested that CL-based, FL-based, and IL-based agents selected more good decks than bad decks, which was more consistent with the subjects' behavior compared to CIL-based agents. There were multiple possible strategies that the models could follow to obtain the final proportions of each deck. Here we aim to examine whether the models were using similar strategies to that used by subjects. We examined the fluctuation of deck preferences of the subjects and model agents over trials. The 95 trials were divided into 10 blocks of 10 trials for the first 9 blocks and 5 trials for the last block. The proportion of each deck selection in each block and the learning scores (i.e. the difference between the number of good deck selections and the number of bad deck selections) for each subject was calculated.

Figure 7.7 shows the learning scores across ten blocks of the IGT. A learning process was apparent both for experimental subjects and CL-based and FL-based agents, in which the learning score progressively improved over blocks, although there was a clear dip in block 10 for subjects and FL-based agents and block 8 for CL-based agents. To quantify the differences between the subjects and the two kinds of model agents over blocks on learning scores, the repeated measures ANOVA tests were conducted in the form of 10 (blocks) \times 2 (groups). The ANOVA test conducted between human and CL-based agents and human and FL-based agents revealed no significant interaction between the two factors (CL-based: $F(7.44, 7290) = 1.39$, $p = 0.20$, FL-based: $F(7.1, 6954) = 1.82$, $p = 0.08$), however, there was a significant main effect of blocks (CL-based: $F(7.44, 7290) = 67.63$, $p < 0.001$, FL-based: $F(7.1, 6954) = 60.64$, $p < 0.001$), with the interpretation that the learning scores of both subjects and CL-based and FL-based model agents varied across




Figure 7.7: The IGT learning scores across ten different blocks of experimental subjects and simulated model agents. The learning process of human subjects, CL-based and FL-based agents was very similar to each other, overall the learning score progressively increasing over blocks. The IL-based agents were least similar to the human subjects, not presenting obvious learning trend in the process. Although CIL-based agents demonstrated similar trends with human subjects, the magnitudes of their learning score were always lower than that of the subjects.

blocks and the CL-based and FL-based model did not behave significantly different across subjects in the learning process. The ANOVA test performed between human and the other two collaborative agents showed that the effects of blocks and group were both significant on the learning score (CIL-based: $F(7.85, 7693) = 47.11, p = 0.03, F(1, 980) = 17.61, p < 0.001$; IL-based: $F(7.71, 7558) = 22.66, p = 0.02, F(1, 980) = 5.69, p = 0.02$) and there was a significant interaction between groups and blocks on the learning score (CIL-based: $F(7.85, 7693) = 5.98, p < 0.001$, IL-based: $F(7.71, 7558) = 15.76, p < 0.001$). This result suggested that the learning process of CIL-based and IL-based agents were significantly different from the experimental subjects, which can also be seen from Figure 7.7, in which the learning scores of CIL-based were lower than the subjects almost over all blocks and the IL-based agents did not present any apparent learning trend in the process.

The Switch Probability after Receiving Losses

Finally, we investigated the immediate effect of losses on subsequent choices. The reason why reward effect was not considered here is subjects could always get rewards no matter which deck they chose, 100 for Deck A and Deck B and 50 for Deck C and Deck D, according to the payoff scheme. Figure 7.8 shows the effect of receiving a loss in the previous trial on the probability of switching to the other deck in the next trial. For the experimental subjects, receiving a loss significantly increased the probability of switching to other decks ($\eta = 0.218, SE = 0.010, p < 0.001$). As the figure shows, same pattern was established by all collaborative agents (CL-based: $\eta = 0.209, SE = 0.007, p < 0.001$; FL-based: $\eta = 0.201, SE = 0.007, p < 0.001$; CIL-based: $\eta = 0.246, SE = 0.007, p < 0.001$; IL-based: $\eta = 0.179, SE = 0.007, p < 0.001$), suggesting that the strategy of loss avoidance to losses used by the four collaborative models was similar to that seen in the subjects' behavior. However, the switch probabilities of CL-based agents were closer to that of the experimental subjects visually and FL-based agents rank second. The switch probability of IL-based agents was clearly higher than the subjects when there was no loss.

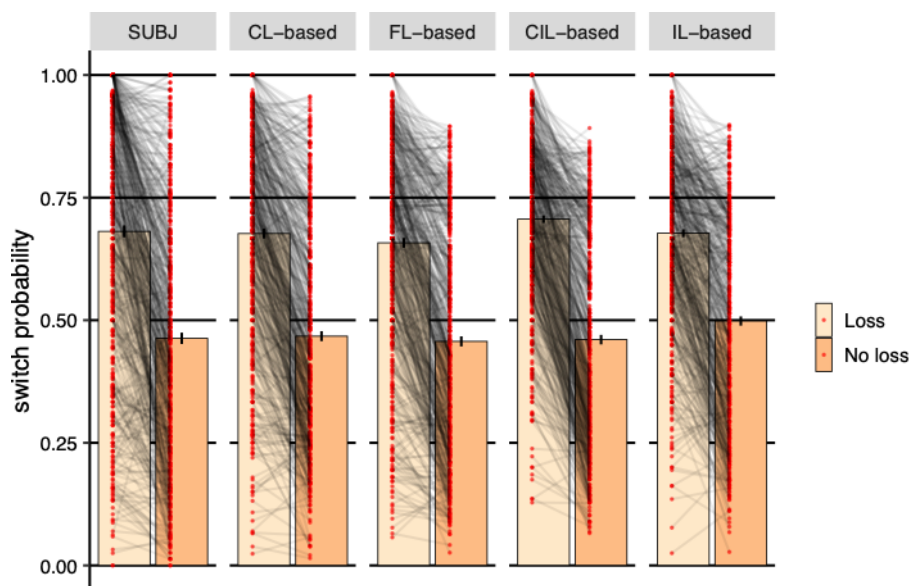


Figure 7.8: Probability of switching to different deck based on whether received losses or not in the previous trial, average over subjects. The CL-based agents demonstrated almost the same switching strategy as human subjects. FL-based agents ranked second. The switch probability of IL-based agents was obviously higher than that of the subjects when there was no loss.

7.3.5 Off-policy Simulation

Off-policy refers to a model which uses previous actions and payoffs to make predictions about the next action. However, the next actions actually used to simulate the models are not derived from these predictions, instead they are derived from the choices made by the human subjects. With such approach we can control the inputs the model receives and examine how they affect the predictions.

Simulations of the models are shown in figure 7.9. Each panel shows a separate simulation across 30 trials (horizontal axis). For the total of 30 trials, the action that was fed into the model was Deck A (top left), Deck B (top right), Deck C (bottom left), or Deck D (bottom right) for each simulation, respectively. The reward and losses associated with these trials were the same as specified in the payoff scheme of IGT. As mentioned earlier, based on the IGT payoff scheme, the player would always receive rewards no matter which deck they chose and only the reward magnitudes varied depending on which deck they chose, while they only received losses occasionally. We only marked trials where there were losses with blue dots in the graphs.

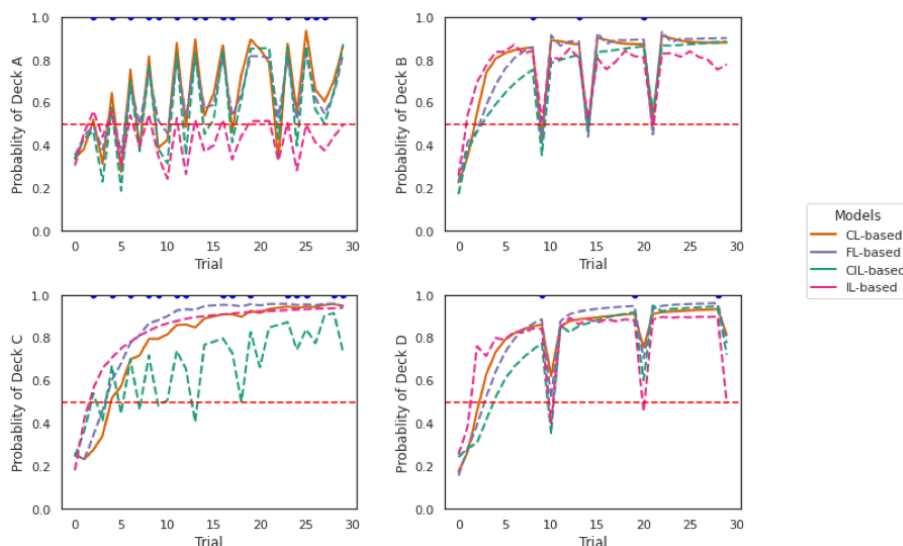


Figure 7.9: Off-policy simulations of the two collaborative models. Each panel shows a simulation of 30 trials. The four collaborative models behaved quite similarly in responding to the losses, especially the FL-based and CL-based model. CIL-based model was more sensitive to the losses given by Deck C as deeper dips are caused compared to other models. IL-based model presents the least perseverance with Deck A compared with other models.

The four collaborative models behaved quite similarly in responding to the losses according to the plots, especially the FL-based model and the CL-based model. Receiving a loss from all decks caused a dip in the probability of staying with that deck, which showed a tendency to switch to another deck. This was consistent with the observations we obtained in Figure 7.8 that the probability of switching was higher when receiving a loss compared to no loss. It seems the CIL-based model was more sensitive to the losses given by Deck C as deeper dips were caused compared to other models. The dips caused by Deck B and Deck D were the deepest, which is not surprising, because Deck B and Deck D are decks with less frequent but larger losses. It is reasonable that the switch probability should be higher if a participant received a significant loss from that deck choice. The probability of choosing Deck A, Deck B, Deck C, and Deck D was higher than that of other decks respectively in each simulation. This implies that persisting with the previously selected deck is a common deck selection tactic for the four models. We can also see that the perseverance with Deck A was not as strong as for the other three decks as the probability of selecting Deck A predicted by the two models was almost never more than 80% and this was even more obvious for the IL-based model. This result was consistent

with the result reported in Figure 7.7 where the proportion of selecting Deck A was the lowest both across subjects and trials. Interestingly, the perseverance effect for Deck C was smaller and the sensitivity to the losses of Deck C was higher for the CIL-based model and relative to other models.

7.4 Discussion

The results demonstrated, as expected, that models trained on a single laboratory's dataset only, performed poorly both on their own testing sets and testing sets from other laboratories. The average prediction accuracy of the model of lab 1, which had the smallest number of subjects among studies, on ten testing datasets was 31%, only 6% higher than the chance correct probability. These results highlighted the need for larger and more representative datasets and therefore emphasized the relevance of training using larger sample size or collaborative models involving subjects from other studies.

Given the difficulties of large-scale data collection and data sharing across studies, institutions and jurisdictions, it is attractive to use distributed learning paradigms to replace traditional centralized training to make large-scale data availability easier. However, the decision to adopt such as an approach is only justifiable so long as the resultant model yields comparable performance with the data pooling approach. While the work presented here cannot be fully comprehensive, it presents for the first time a representative sampling of eight collaborative models, one of which was based on centralized data training while the other seven were based on distributed learning including one FL-based, three CIL-based and three IL-based, were trained on the training datasets from all laboratories for comparison. It was found that the order of the laboratories was an important factor that influenced the performance of CIL-based and IL-based models. An ascending order based on the laboratory dataset sizes was the best strategy for both of these two paradigms and the corresponding models were selected, along with the FL-based model, to be further evaluated with the centralized model from multiple aspects. When examining their prediction capability on the unseen testing datasets the prediction accuracy was improved by 24%, 24%, 24%, and 23% points compared to the worst single model on

the same testing data-sets for CL-based, FL-based model, CIL-based and IL-based model respectively. This highlighted the value of involving larger numbers of data samples for training a deep learning model. It is also clear that FL-based and CIL-based approaches can both achieve competitive performance with CL-based model while retaining the benefit of enhanced data privacy. Another interesting finding was that all the collaborative models performed clearly better on the testing sets of lab 9, lab 5, and lab 4 and the CIL-based and IL-based models did not obviously suffer from the 'catastrophic forgetting' problem as reported in the literature [252]–[254]. This is apparent because they did not present performance biased to the data they had mostly recently seen (i.e. lab 2 for the ascending order case, lab 1 in the descending order case, and lab 10 for the random order case).

We then evaluated the capability of the distributed model in capturing characteristics of human decision-making strategies used on the IGT in comparison with CL-based model through on-policy and off-policy simulations. In the on-policy simulation, 491 fake agents, which was the same as size as the training data-set for training the collaborative models, were generated for the four models. The results revealed that CL-based agents behaved more similarly to human subjects exhibiting similar choice proportions on all decks, while the agents generated by the three distributed models had significantly different deck preferences and aversions on one or two particular decks compared to the subjects' choices. This characteristic was also reflected in the fluctuation of the probability of selecting each deck over trials except that FL-based agents performed as well as CL-based agents in this comparison. The comparison analysis of the IGT learning scores across trials between the experimental subjects and model agents also suggested CIL-based and IL-based model were weaker in recognizing the good decks as the learning scores of these two agents did not change for IL-based agents or were always lower than the subjects throughout the task, although the overall trend of the learning score was increasing for CIL-based agents. In terms of the strategy of dealing with losses, all model agents adopted the same strategy as human subjects, i.e. switching more often when they received a loss compared to no loss. It is worth noting here that although no statistically significant weakness was identified for CIL-based and IL-based agents, the average switching probability of receiving losses

was relatively higher than that of the human subjects visually, while the CL-based and FL-based agents were almost the same as the subjects on this figure. The characteristic of switching more after losses was validated again in the following off-policy simulation, where the actions fed to the model were specified by us and the models were responsible for predicting the next action to be taken for the next trial. According to the model's prediction, the probability of sticking with a deck was decreased whenever there a loss incurred. CL-based and FL-based models reacted almost the same way again in this simulation, but IL-based model demonstrated less perseverance to Deck A and the CIL-based model was more sensitive to losses compared to the other agents.

In summary, consistent with the results where FL demonstrated comparative performance on different benchmark datasets, [262]–[265], the FL-based model in this paper achieved the best prediction accuracy and capability in learning and mimicking the decision-making strategies used by the experimental subjects among the other distributed learning paradigms. Although CIL-based and IL-based models also achieved considerable prediction accuracy when compared to the CL-based model, they failed in capturing some features such as subtle deck preferences between Deck B and Deck C and learning curve shapes for the subjects. Although in this study the distributed learning strategies were applied to a simulated scenario where data from different laboratories was isolated because of barriers of data sharing regulations, it has implications to deal with challenges encountered in many situations for formulating large-scale datasets in the field of cognitive and behavioral science. For example, if the researcher was prevented from collecting a much wider range of private passive data from subjects' smartphones due to data privacy and protection policies, the distributed learning strategies examined in this paper could be an alternative approach to use for them. It is worth speculating what deeper insights an interpretation of the parameters might reveal in terms of human choices under uncertainty. If important information describing underlying human computational elements is encoded through subtle variations in the parameter values, then the noise or uncertainty introduced through distributed learning needs to be carefully scrutinized. It suggests that simulated datasets may be a useful means for determining to what extent distributed learning com-

promises the quality of the models developed. Importantly, the comparisons of methods introduced in this paper is informative regarding an approach to select the best technique. Additionally, emerging improvements to FL [266], [267] hold further promise for the narrowing of the gap in performance between centralized and federated learning.

7.5 Conclusion

In this study, the feasibility and reliability of applying distributed learning paradigms for characterizing and predicting human decision-making in scenarios where large-scale data formulating is hindered by ethical and law regulations. It was found that FL outperformed the other distributed learning paradigms in learning the learning features of human subjects in this context. The degradation in performance of the FL-based model compared to the CL-based model in terms of the capabilities of predicting and learning human behavior was not negligible and its impact depends on the subsequent use of the model. Nonetheless, the use of FL-based over CL-based model has the immediate advantage of keeping data confidential, which is an important requirement enshrined in policy and law. An understanding of the trade-offs inherent in data privacy versus accuracy of measurement is important and this paper demonstrates an approach for assessing this in the context of a concrete, real-world multi-centre data collection effort. We hope that the development of these techniques and emerging derivations of the approach will allow for data-private distributed training over large-scale subjects with various ages, education backgrounds and social status. It will facilitate the development of robust and reliable digital biomarkers in the context of computational phenotyping of human decision-making and ultimately maximum translation from basic research to clinic.

Chapter 8

Conclusion

This chapter will conclude the thesis and reflect on the key research findings in relation to the research objectives and questions and discussing the value and contribution thereof. It will also review the limitations of the study and propose opportunities for future research.

8.1 Summary

The infusion of theories from economics and artificial intelligence into neuroscience has dramatically advanced our understanding of the underlying neural mechanisms for decision-making and reinforcement learning with conventional lab-based hard neuroscience research [2], [5], [11], [268]. Nevertheless, improvements in understanding neurobiological mechanisms of mental health have not been translated into clinical settings to solve clinical problems. In addition, there is limited consensus on the specific barriers that hinder the progress toward translation, as well as the best strategies for overcoming these barriers. One of the many factors that impedes the clinical translation is that traditional single-site and one-shot laboratories studies with small sample size are not sufficient for capturing numerous interacting and confounding variables within the same individual and thus developing strong neurocognitive models of mental disorders. New methodologies that can deliver on much broader and multidimensional data sets are needed to bridge the gap between basic cognitive research to clinical application. In recent years, the technological advances in data collection methods, e.g. the advent of smartphones and

sensor-based wearables, have presented opportunities to overcome some of the limitations inherent to conventional laboratory experiments [26], [27], [56]. This thesis contributes to research both experimentally and methodologically that has potential to advance the clinical translation of computational phenotyping of human decision-making. There are three aspects to these contributions: moving the experimental site outside the lab, capturing the temporal fluctuations in cognitive function and clinical symptoms, and overcoming the ethical and legal barriers of large-scale phenotyping. We proposed to implement the two methodologies, lab-in-the-field [19] and EMA experiments [20], in behavioral and cognitive research to capture how behavioral phenotypes manifest in the real world, change over time, and interact with emotional, physical and social states. We also proposed to use distributed learning strategies to overcome the ethical and legal restrictions on large-scale human phenotyping research. Lower cost and easier accessibility to large-scale phenotyping study will promote the development of robust and reliable cognitive models and ultimately will be beneficial to clinical translation. The proposed experimental and data training strategies were examined in four case studies to demonstrate the feasibility and reliability of implementing them into the basic research.

In the first case study in **Chapter 4**, the lab-in-the-field method was implemented in a study where we aimed to capture differences in the decision-making phenotypes on the IGT between patients who suffer from chronic pain and healthy controls in naturalistic settings. Significant differences were identified in person-specific computational phenotypes that correspond to key theoretical mechanisms when comparing chronic pain patients and healthy controls. Chronic pain individuals were more sensitive to rewards and less persistent in their choices during the IGT compared to controls and the altered computational parameters could predict subjects' pain severity and pain interference. In the second case study in **Chapter 5**, the lab-in-the-field method was applied to patients with PD. The lab-in-the-field method is particularly valuable for such specific group of population who may have difficulties in traveling and remote assessment is a better option. Using advanced computational modeling techniques, we found that the deficits in adaptation of learning rates to contingency volatility on the ALT of PD patients were associated the

depression and anxiety symptoms in this population. The successful implementation of the two lab-in-the-field studies and the significant findings of the variables of interest in the noisier naturalistic environments demonstrated it is feasible to move the behavioral paradigms outside the laboratory and to characterizing human decision-making in their daily life contexts even for vulnerable groups such as patients with PD. In fact, switching experimental data collection from the laboratory to subject's everyday living contexts has been a trend, particularly in the field of psychiatric research [55], [269]. Although it is still relatively new, uptaking of smartphones for research, or more general internet-based phenotyping of human decision-making, have generated new insights into psychiatrically relevant phenomena [198], [270]–[272].

The third case study in **Chapter 6** applied the EMA methodology to patients suffering from chronic tinnitus, a disease known to have fluctuations in symptoms related to perceived distress. Participants in this study were sampled frequently for their momentary well-being experiences combined with cognitive testing assessed by the ALT at various time points and situations using an EMA app. It was observed that the fluctuation of the tinnitus stressfulness was associated with performance variation of adaptive learning to the volatility of the environment. In addition to observe how computational phenotypes of decision-making manifest in the real world, the EMA methodology allows us to embrace the dynamic process inherent to clinical symptoms and potential correspondence between cognitive function and clinical symptoms. Specifically, the application of a mobile app in this study facilitated closer temporal monitoring of symptoms than is typically not feasible for traditional laboratories studies. The method of sampling contemporary experiences and triggering cognitive tests of decision-making in this case study is applicable to many other psychological research that aim at providing insights in the causal relationships that exist between cognitive abilities and within-individual factors such as sleep [273], alcohol [274], and drug use [275], etc. Combining with the computational parsing of the latent drivers of learning and decision-making, such EMA studies has potential to identify which cognitive decision processes might reflect changes at short timescales (e.g., days, weeks, or months) that can guide ongoing treatment or prognosis, and which might primarily reflect stable

features of the disorder that could inform early prevention efforts.

The last case study in **Chapter 7** presented the application of the distributed learning strategies to fit DL models to characterize decision-making without access to the raw data. The comparable performance achieved by the FL strategy compared to the straightforward centralized pooling method indicated that the difficulties of large-scale human data availability caused by the ethical and law regulations could be addressed by advanced privacy-preserving AI techniques. The FL strategy was originally designed for protecting data on mobile devices, thus it is particularly suitable for smartphone-based research methods advocated in this thesis. Removing the barriers of incorporating large-scale data in model training is of significant importance for scaling up the range of the investigated individuals, thus substantially increasing the representativeness and diversity of the samples. The findings in this case study suggest that FL might be an easier and more practical way of implementing large-scale phenotyping in the future to overcome the stricter regulations to data and personal privacy protection.

In summary, this thesis proposed to use new research methodologies in behavioral and cognitive research to further advance our understanding of cognitive mechanisms and accelerate the clinical translation and successfully demonstrated their feasibility for a suite of new investigations with the implementation of the approaches in case studies. The attempt to dramatically extend the depth and breath of computational phenotyping research using the proposed methodologies is just in its infancy, but the successful implementations of these methodologies and findings in the case studies provide empirical evidence for the possibility and feasibility of collecting richer, multivariate human decision-making datasets under naturalistic settings among larger samples taking the advantage of smartphones and state-of-the-art machine learning techniques. The findings also contribute to our understanding of the characteristics of decision-making of various populations (chronic pain, PD, and tinnitus) and the influences of potential covariates particularly in settings where achieving the desired internal validity of the measurements is more challenging. It is expected that the smartphone-based methodologies verified in this thesis and their combination with computational tools will be an inexpensive add-on

to a basic cognitive study that can complement the conventional lab-based experiments. The use and further development of the lab-in-the-field and EMA methodologies and the distributed learning strategy may reshape behavioral and cognitive research and clinical trials in the near future.

8.2 Limitations

It is acknowledged by the author that the applications of these methodologies in the case studies were not without limitations and challenges. In the chronic pain study, only gender, age, and pain duration were collected in the demographic data, but the mental health conditions of the participants were not considered apart from pain severity and pain interference. It has been known that experiencing chronic pain puts a person at increased risk for developing anxiety and depression disorders and these two factors have been documented to have significant influences on decision-making performance on the IGT [276], [277]. The impact of chronic pain medication, which may bias participants' response to affective information, was not taken into consideration either. It is also worth noting that virtual money may not be as effective as real financial incentives for the decision-making task. However, we did get good adherence and compliance for the experiment. Furthermore, we recruited a mixed sample of subjects who might have had various kinds of pain conditions, which might cause various cognitive abnormalities and thus have different influences on decision-making [145]. As a result, caution must be exercised when interpreting the results obtained and generalizing them to pain populations in specific environments and particularly chronic pain conditions. For the PD study, one of the limitations was that the correlational study analysis design may limit the ability to draw cause and effect explanations [188], which is also the case for the tinnitus study where no healthy control group was recruited. In future studies, we will incorporate PD patients with and without depression and anxiety and compare their performance on the adaptation of learning to volatility in order to obtain concrete causal conclusions. In the tinnitus study, regardless the great effort we have made on the recruitment process, the sample size ended up being relatively small. Assessments such as the behavioral paradigms

that require a substantial amount of time to complete cannot be easily incorporated into participants' daily routines or administered frequently without overly inconveniencing users. In particular the fixed thresholds for triggering the task used in the study was very likely to cause an excessive number of triggers. I have contributed to the design of an adaptive triggering algorithm that has been designed to minimise participant burden [278]. However, the algorithm had not been embedded into the functionality of AthenaCX at the time of the study. Thus, there is a long way to go in terms of the optimal experimental design in the case of dense sampled EMA.

From the global perspective of the thesis, the samples sizes of the case studies conducted were all relatively small, especially given the less controlled experimental methodologies adopted in this thesis. It would be more comprehensive to present a case study in which a large-scale lab-in-the-field or EMA experiment using smartphones is performed to illustrate the feasibility of doing so. Due to the time and cost limit, we ended with providing a solution for addressing the ethical and legal issues encountered on the way of implementing large-scale phenotyping. We take the smaller scale of studies as pilot research mainly for validating our research methodologies, i.e. the possibilities of phenotyping characteristics of human decision-making with computational models and capturing their variations in naturalistic settings. In fact, this is absolutely not the last issue we have on the pathway of realizing the goals of computation to clinical translation. As the complexity of the experiments increases, the logistic challenges to correctly deliver them also increase and efficiency in data collection becomes necessary to reduce the costs of a prolonged experiment. As mentioned earlier, in addition to active data like repeated cognitive testing and experience sampling, sensors on smartphones can provide additional complementary tools for collecting numerous passive data, which can be used to create so-called digital phenotypes [279]. This type of data has not been included in this thesis. Apart from the ethical and legal issues, another difficulty exists in analysing these data. Passive data is generally complex and noisy with countless features. The addition of features can present challenges for robust analysis, and as smartphone-based data sets become larger, this problem can get worse, and external validation becomes increasingly important. Although

ML approaches and dimensionality-reduction methods have been adopted to cope with the curse of dimensionality, combining passive data with active data and deriving additional insights into the underlying mechanisms is still not a easy task. In addition, a limitation of the smartphone-based methodologies examined in this thesis is that they can not be used to simultaneously collect some of the hard measures since fMRI and positron emission tomography scans are unlikely to ever be collected remotely. This may pose challenges in fully comprehending certain cognitive processes. However, investment on large-scale smartphone methodologies does not mean complemently abandoning conventional lab-based experiments. Instead, they should be complementary to each other, the former gathering vast but noisy data, while the latter gathering smaller data sets of higher detail and quality.

8.3 Future Work

8.3.1 Further Applications of the Lab-in-the-Field Methodology

Apart from the aforementioned lab-in-the-field studies involving chronic pain and PD patients, we have also applied this methodology to a number of challenging populations as part of efforts to further support the utility of the approaches being advocated here. We begin with an account of efforts to engage with a particularly challenging group - people who have strong beliefs in conspiracy theories. Belief in conspiracy theories has been prevalent throughout human history at various times and in various societies [280]–[282]. The outbreak of COVID-19 further accelerated the spreading and development of conspiracy theories. Although conspiracy theories vary enormously in content, it has been noted that belief in such theories is grounded in similar underlying psychological processes. It has been found that numerous psychological traits are associated with belief in conspiracy theories, including distrust toward authorities, high motives to conspire, feelings of out of control, feelings of dissatisfaction about one’s life, etc [283]–[287]. The emotional state of anxiety was also linked to a higher belief in conspiracy theories [288], [289], as was the political attitude of right-wing authoritarianism [290]. The studies

that focused on examining the role of individual differences in cognitive mechanisms of decision-making are rarely seen in the literature. It was found people who prone to believe in irrational theories are more desired to seek for existing patterns to establish causal relations even in facing with a random process where in fact no pattern exists, but they were less able to recognize the real patterns when they are presented with a nonrandom process where a pattern did exist [291]. In the pilot study which examines 39 individuals, we explored if the general tendency to engage with conspiracy theories was associated with deficits in recognizing and learning the probabilistic patterns of the environment in the ALT. The participants were recruited online and were able to enter into the experiment page directly through a URL attached. There were required to rate their opinions on The Generic Conspiracist Beliefs Scale and then complete the ALT there. It was found that high conspiracy belief individuals had difficulties differentially updating action-outcome contingencies as a function of whether the current environment was stable or volatile in the ALT. The difficulty in identifying the stability of action-outcome contingencies to judge whether or not to repeat an action that has led to an unexpected aversive outcome may lead high conspiracy individuals to engage in poor decision-making. It might also result in aversive outcomes being experienced as less predictable and less avoidable. This could in turn increase in desire to make sense of the environment, and thus potentially be involved in the onset of conspiracist ideation. In the future, more comprehensive variables will be evaluated to explain the relationship between impaired pattern recognition of the probabilistic structure of the environment and conspiracy belief.

8.3.2 Extension of the EMA Method

Of clinical importance is whether the alterations of the computational phenotypes that characterize underlying cognitive processes of decision-making are sensitive to treatment, which would indicate a potential causal relationship between phenotype changes and symptom improvement. While carefully monitored, it is simply difficult to follow enough patients for enough time period to observe the changes of clinical variables in sufficient numbers for meaningful analysis with traditional lab-based treatment studies. Another

source of variability of decision-making is emotion. Although several techniques for inducing emotion responses (e.g. induction of stress) in the laboratory have been developed, observing daily fluctuations in mood and decision-making in naturalistic settings is crucial to delineate the complex means by which affect can influence choices. The experimental methodologies, particularly the EMA method, advocated and examined in this thesis are poised to achieve the goal of convenient longitudinal tracking of symptoms and behaviors in individuals. They are of great potential to provide a theoretically rich, predictive, and detailed science of decision-making. Simultaneously, the distributed learning strategy will be useful in scenarios, in which the subjects are recruited from multiple different clinical centers where the treatments are actually prescribed and there are ethical issues for sharing the patients' data among the collaborators.

8.3.3 Scaling the Experiment with the Voice-based Devices

Another future direction concerns scaling the computational phenotyping of human decision-making through the pervasive voice-based computing devices to create larger datasets. Decision-making research has focused for a long time on investigating reinforcement learning and decision-making from visual input or symbolic information. However, people receive information and interact with their environments through various senses. In particular, our sense of hearing in conjunction with speech represents a ubiquitous modality for learning and for updating our knowledge of the world. Thus, in the future, we aim to expand the experimental approaches beyond conventional visual representations. This represents an important path for the investigation of human decision-making which is now experimentally accessible via rapid advances in voice-enabled intelligent personal assistants (IPAs). Examples include Amazon's Alexa technology and Google's Voice Assistant. However, to date no studies have demonstrated the feasibility of delivering such experimental paradigms over such voice technologies. We have conducted a preliminary study in which the performance of the same group of participants ($N = 30$) on the traditional visual-based the conversational voice-based two-armed bandit task were compared. Both model-independent behavioral measures and model-derived parameters

were compared. The results suggested that participants demonstrated higher shifting rates in the voice-based version of the task. The computational modelling analysis revealed that participants adopted similar learning rates under the two versions of the interfaces, but more decision noise was introduced in the voice-based task as reflected by the decreased value of the inverse temperature parameter. We suggest that the elevated shifting rate is derived from the increased noise in the voice interface instead of a change in the learning strategy of the participants. Higher intensity of the control adjustments (click touch versus speak) might be one of the sources of noise, thus it is important to think further regarding the design of the voice interface if we wish to apply voice-enabled IPAs to measure human decision-making in their daily environments in the future. A publication detailing this work is available.

8.3.4 Summary

Pervasive computing devices like smartphones and IPAs have changed our lives in unprecedented ways. One benefit of these ubiquitous technologies is their ability to create opportunities for the development of better tools to aid in behavioral and cognitive studies. The continuing and rapid development of computational tools holds great promise for interpreting the data acquired with these devices and delineating the underlying mechanisms by which people adapt to their environments. This includes those that lead to maladaptive forms of learning and decision-making within clinical populations. It is within this context that the contributions of this thesis are best understood and through these contributions we hope we can further advance the positive contribution of machine learning to help people better understand themselves.

Bibliography

- [1] A. G. Collins and A. Shenhav, “Advances in modeling learning and decision-making in neuroscience,” *Neuropsychopharmacology*, vol. 47, no. 1, pp. 104–118, 2022.
- [2] P. W. Glimcher and E. Fehr, *Neuroeconomics: Decision making and the brain*. Academic Press, 2013.
- [3] A. Tversky and D. Kahneman, “The framing of decisions and the psychology of choice,” in *Behavioral Decision Making*, Springer, 1985, pp. 25–41.
- [4] J. W. Kable and P. W. Glimcher, “The neurobiology of decision: Consensus and controversy,” *Neuron*, vol. 63, no. 6, pp. 733–745, 2009.
- [5] D. Lee, H. Seo, and M. W. Jung, “Neural basis of reinforcement learning and decision making,” *Annual Review of Neuroscience*, vol. 35, p. 287, 2012.
- [6] L. A. Zadeh, “Outline of a new approach to the analysis of complex systems and decision processes,” *IEEE Transactions on Systems, Man, and Cybernetics*, no. 1, pp. 28–44, 1973.
- [7] P. R. Montague, R. J. Dolan, K. J. Friston, and P. Dayan, “Computational psychiatry,” *Trends in Cognitive Sciences*, vol. 16, no. 1, pp. 72–80, 2012.
- [8] K. T. Kishida, B. King-Casas, and P. R. Montague, “Neuroeconomic approaches to mental disorders,” *Neuron*, vol. 67, no. 4, pp. 543–554, 2010.
- [9] A. D. Redish, *The mind within the brain: How we make decisions and how those decisions go wrong*. Oxford University Press, 2013.

- [10] R. A. Adams, Q. J. Huys, and J. P. Roiser, “Computational psychiatry: Towards a mathematically informed understanding of mental illness,” *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 87, no. 1, pp. 53–63, 2016.
- [11] Q. J. Huys, M. Browning, M. P. Paulus, and M. J. Frank, “Advances in the computational understanding of mental illness,” *Neuropsychopharmacology*, vol. 46, no. 1, pp. 3–19, 2021.
- [12] K. J. Friston, A. D. Redish, and J. A. Gordon, “Computational nosology and precision psychiatry,” *Computational Psychiatry (Cambridge, Mass.)*, vol. 1, p. 2, 2017.
- [13] C. Gonzalez, “Decision-making: A cognitive science perspective,” *The Oxford Handbook of Cognitive Science*, p. 249, 2016.
- [14] K. Danziger, *Constructing the subject: Historical origins of psychological research*. Cambridge University Press, 1994.
- [15] G. Hatfield, “Psychology, philosophy, and cognitive science: Reflections on the history and philosophy of experimental psychology,” *Mind & Language*, vol. 17, no. 3, pp. 207–232, 2002.
- [16] U. Bronfenbrenner, “Toward an experimental ecology of human development.” *American Psychologist*, vol. 32, no. 7, p. 513, 1977.
- [17] A. Kingstone, D. Smilek, and J. D. Eastwood, “Cognitive ethology: A new approach for studying human cognition,” *British Journal of Psychology*, vol. 99, no. 3, pp. 317–340, 2008.
- [18] S. G. Shamy-Tsoory and A. Mendelsohn, “Real-life neuroscience: An ecological approach to brain and behavior research,” *Perspectives on Psychological Science*, vol. 14, no. 5, pp. 841–859, 2019.
- [19] U. Gneezy and A. Imas, “Lab in the field: Measuring preferences in the wild,” in *Handbook of Economic Field Experiments*, vol. 1, Elsevier, 2017, pp. 439–464.
- [20] S. Shiffman, A. A. Stone, and M. R. Hufford, “Ecological momentary assessment,” *Annu. Rev. Clin. Psychol.*, vol. 4, pp. 1–32, 2008.

- [21] G. W. Harrison and J. A. List, “Field experiments,” *Journal of Economic Literature*, vol. 42, no. 4, pp. 1009–1055, 2004.
- [22] J. A. List, “Field experiments: A bridge between lab and naturally occurring data,” *The BE Journal of Economic Analysis & Policy*, vol. 6, no. 2, 2007.
- [23] J. Henrich, S. J. Heine, and A. Norenzayan, “The weirdest people in the world?” *Behavioral and Brain Sciences*, vol. 33, no. 2-3, pp. 61–83, 2010.
- [24] P. Voigt and A. Von dem Bussche, “The eu general data protection regulation (gdpr),” *A Practical Guide, 1st Ed., Cham: Springer International Publishing*, vol. 10, p. 3 152 676, 2017.
- [25] C. Gwaltney, S. J. Coons, P. O’Donohoe, *et al.*, ““bring your own device”(byod): The future of field-based patient-reported outcome data collection in clinical trials?” *Therapeutic Innovation & Regulatory Science*, vol. 49, no. 6, pp. 783–791, 2015.
- [26] A. Seifert, M. Hofer, and M. Allemand, “Mobile data collection: Smart, but not (yet) smart enough,” *Frontiers in Neuroscience*, vol. 12, p. 971, 2018.
- [27] G. M. Harari, N. D. Lane, R. Wang, B. S. Crosier, A. T. Campbell, and S. D. Gosling, “Using smartphones to collect behavioral data in psychological science: Opportunities, practical considerations, and challenges,” *Perspectives on Psychological Science*, vol. 11, no. 6, pp. 838–854, 2016.
- [28] S. D. Gosling and W. Mason, “Internet research in psychology,” *Annual Review of Psychology*, vol. 66, pp. 877–902, 2015.
- [29] M. Webster and J. Sell, *Laboratory experiments in the social sciences*. Elsevier, 2014.
- [30] E. H. Patzelt, C. A. Hartley, and S. J. Gershman, “Computational phenotyping: Using models to understand individual differences in personality, development, and mental illness,” *Personality Neuroscience*, vol. 1, 2018.

- [31] B. E. Roe and D. R. Just, "Internal and external validity in economics research: Tradeoffs between experiments, field experiments, natural experiments, and field data," *American Journal of Agricultural Economics*, vol. 91, no. 5, pp. 1266–1271, 2009.
- [32] A. A. Stone and S. Shiffman, "Ecological momentary assessment (ema) in behavioral medicine.," *Annals of Behavioral Medicine*, 1994.
- [33] S. Shiffman, "Dynamic influences on smoking relapse process," *Journal of Personality*, vol. 73, no. 6, pp. 1715–1748, 2005.
- [34] M. Csikszentmihalyi *et al.*, "The ecology of adolescent activity and experience.," *Journal of Youth and Adolescence*, vol. 6, no. 3, pp. 281–94, 1977.
- [35] M. Conner, M. Fitter, and W. Fletcher, "Stress and snacking: A diary study of daily hassles and between-meal snacking," *Psychology and Health*, vol. 14, no. 1, pp. 51–63, 1999.
- [36] S. Shiffman, M. H. Balabanis, C. J. Gwaltney, *et al.*, "Prediction of lapse from associations between smoking and situational antecedents assessed by ecological momentary assessment," *Drug and Alcohol Dependence*, vol. 91, no. 2-3, pp. 159–168, 2007.
- [37] C. J. Burgin, P. J. Silvia, K. M. Eddington, and T. R. Kwapil, "Palm or cell? comparing personal digital assistants and cell phones for experience sampling research," *Social Science Computer Review*, vol. 31, no. 2, pp. 244–251, 2013.
- [38] S. T. Doherty, C. J. Lemieux, and C. Canally, "Tracking human activity and well-being in natural environments using wearable sensors and experience sampling," *Social Science & Medicine*, vol. 106, pp. 83–92, 2014.
- [39] W. Hofmann and P. V. Patel, "Surveysignal: A convenient solution for experience sampling research using participants' own smartphones," *Social Science Computer Review*, vol. 33, no. 2, pp. 235–253, 2015.

- [40] P. C. Clasen, A. J. Fisher, and C. G. Beevers, "Mood-reactive self-esteem and depression vulnerability: Person-specific symptom dynamics via smart phone assessment," *PLoS One*, vol. 10, no. 7, e0129774, 2015.
- [41] P. Klasnja and W. Pratt, "Healthcare in the pocket: Mapping the space of mobile-phone health interventions," *Journal of Biomedical Informatics*, vol. 45, no. 1, pp. 184–198, 2012.
- [42] J. Ainsworth, J. E. Palmier-Claus, M. Machin, *et al.*, "A comparison of two delivery modalities of a mobile phone-based assessment for serious mental illness: Native smartphone application vs text-messaging only implementations," *Journal of Medical Internet Research*, vol. 15, no. 4, e2328, 2013.
- [43] R. Band, C. Barrowclough, R. Emsley, M. Machin, and A. J. Wearden, "Significant other behavioural responses and patient chronic fatigue syndrome symptom fluctuations in the context of daily life: An experience sampling study," *British Journal of Health Psychology*, vol. 21, no. 3, pp. 499–514, 2016.
- [44] P. F. Cook, S. J. Schmiege, L. Bradley-Springer, W. Starr, and J. M. Carrington, "Motivation as a mechanism for daily experiences' effects on hiv medication adherence," *Journal of the Association of Nurses in AIDS Care*, vol. 29, no. 3, pp. 383–393, 2018.
- [45] L. J. Faherty, L. Hantsoo, D. Appleby, M. D. Sammel, I. M. Bennett, and D. J. Wiebe, "Movement patterns in women at risk for perinatal depression: Use of a mood-monitoring mobile application in pregnancy," *Journal of the American Medical Informatics Association*, vol. 24, no. 4, pp. 746–753, 2017.
- [46] M. Fuller-Tyszkiewicz, T. Karvounis, R. Pemberton, L. Hartley-Clark, and B. Richardson, "Determinants of depressive mood states in everyday life: An experience sampling study," *Motivation and Emotion*, vol. 41, no. 4, pp. 510–521, 2017.

- [47] S. D. Gosling, S. Vazire, S. Srivastava, and O. P. John, "Should we trust web-based studies? a comparative analysis of six preconceptions about internet questionnaires.," *American Psychologist*, vol. 59, no. 2, p. 93, 2004.
- [48] A. N. Joinson, A. Woodley, and U.-D. Reips, "Personalization, authentication and self-disclosure in self-administered internet surveys," *Computers in Human Behavior*, vol. 23, no. 1, pp. 275–285, 2007.
- [49] L. Germine, K. Nakayama, B. C. Duchaine, C. F. Chabris, G. Chatterjee, and J. B. Wilmer, "Is the web as good as the lab? comparable performance from web and lab in cognitive/perceptual experiments," *Psychonomic Bulletin & Review*, vol. 19, no. 5, pp. 847–857, 2012.
- [50] M. J. Crump, J. V. McDonnell, and T. M. Gureckis, "Evaluating amazon's mechanical turk as a tool for experimental behavioral research," *PloS One*, vol. 8, no. 3, e57410, 2013.
- [51] J. K. Goodman, C. E. Cryder, and A. Cheema, "Data collection in a flat world: The strengths and weaknesses of mechanical turk samples," *Journal of Behavioral Decision Making*, vol. 26, no. 3, pp. 213–224, 2013.
- [52] S. Palan and C. Schitter, "Prolific. ac—a subject pool for online experiments," *Journal of Behavioral and Experimental Finance*, vol. 17, pp. 22–27, 2018.
- [53] J. Chandler and D. Shapiro, "Conducting clinical research using crowdsourced convenience samples," *Annual Review of Clinical Psychology*, vol. 12, pp. 53–81, 2016.
- [54] N. Stewart, J. Chandler, and G. Paolacci, "Crowdsourcing samples in cognitive science," *Trends in Cognitive Sciences*, vol. 21, no. 10, pp. 736–748, 2017.
- [55] C. M. Gillan and R. B. Rutledge, "Smartphones and the neuroscience of mental health," *Annual Review of Neuroscience*, vol. 44, p. 129, 2021.
- [56] H. R. Brown, P. Zeidman, P. Smittenaar, *et al.*, "Crowdsourcing for cognitive science—the utility of smartphones," *PloS One*, vol. 9, no. 7, e100662, 2014.

- [57] R. B. Rutledge, N. Skandali, P. Dayan, and R. J. Dolan, “A computational and neural model of momentary subjective well-being,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 33, pp. 12 252–12 257, 2014.
- [58] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT Press, 2018.
- [59] X.-J. Wang, “Decision making in recurrent neuronal circuits,” *Neuron*, vol. 60, no. 2, pp. 215–234, 2008.
- [60] P. N. Tobler and E. U. Weber, “Valuation for risky and uncertain choices,” in *Neuroeconomics*, Elsevier, 2014, pp. 149–172.
- [61] A. Maurice and O. Hagen, “Expected utility hypothesis and the allais paradox,” *Dordrecht: Reidel*, pp. 27–145, 1979.
- [62] A. Tversky and D. Kahneman, “Advances in prospect theory: Cumulative representation of uncertainty,” *Journal of Risk and Uncertainty*, vol. 5, no. 4, pp. 297–323, 1992.
- [63] J. C. Hershey and P. J. Schoemaker, “Prospect theory’s reflection hypothesis: A critical examination,” *Organizational Behavior and Human performance*, vol. 25, no. 3, pp. 395–418, 1980.
- [64] D. Kahneman and A. Tversky, “Prospect theory: An analysis of decision under risk,” in *Handbook of the Fundamentals of Financial Decision Making: Part I*, World Scientific, 2013, pp. 99–127.
- [65] D. Kahneman and A. Tversky, “Choices, values, and frames,” in *Handbook of the Fundamentals of Financial Decision Making: Part I*, World Scientific, 2013, pp. 269–278.
- [66] P. K. Lattimore, J. R. Baker, and A. D. Witte, “The influence of probability on risky choice: A parametric examination,” *Journal of Economic Behavior & Organization*, vol. 17, no. 3, pp. 377–400, 1992.

- [67] T. A. Hare, C. F. Camerer, and A. Rangel, "Self-control in decision-making involves modulation of the vmPFC valuation system," *Science*, vol. 324, no. 5927, pp. 646–648, 2009.
- [68] S. M. Tom, C. R. Fox, C. Trepel, and R. A. Poldrack, "The neural basis of loss aversion in decision-making under risk," *Science*, vol. 315, no. 5811, pp. 515–518, 2007.
- [69] T. A. Hare, C. F. Camerer, D. T. Knöpfle, J. P. O'Doherty, and A. Rangel, "Value computations in ventral medial prefrontal cortex during charitable decision making incorporate input from regions involved in social cognition," *Journal of Neuroscience*, vol. 30, no. 2, pp. 583–590, 2010.
- [70] J. Peters and C. Büchel, "Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making," *Journal of Neuroscience*, vol. 29, no. 50, pp. 15 727–15 734, 2009.
- [71] S.-W. Wu, M. R. Delgado, and L. T. Maloney, "The neural correlates of subjective utility of monetary outcome and probability weight in economic and in motor decision under risk," *Journal of Neuroscience*, vol. 31, no. 24, pp. 8822–8831, 2011.
- [72] I. Levy, J. Snell, A. J. Nelson, A. Rustichini, and P. W. Glimcher, "Neural representation of subjective value under risk and ambiguity," *Journal of Neurophysiology*, vol. 103, no. 2, pp. 1036–1047, 2010.
- [73] P. N. Tobler, C. D. Fiorillo, and W. Schultz, "Adaptive coding of reward value by dopamine neurons," *Science*, vol. 307, no. 5715, pp. 1642–1645, 2005.
- [74] C. D. Fiorillo, P. N. Tobler, and W. Schultz, "Discrete coding of reward probability and uncertainty by dopamine neurons," *Science*, vol. 299, no. 5614, pp. 1898–1902, 2003.
- [75] S. W. Kennerley, A. F. Dahmubed, A. H. Lara, and J. D. Wallis, "Neurons in the frontal lobe encode the value of multiple decision variables," *Journal of Cognitive Neuroscience*, vol. 21, no. 6, pp. 1162–1178, 2009.

- [76] S. W. Kennerley, T. E. Behrens, and J. D. Wallis, “Double dissociation of value computations in orbitofrontal and anterior cingulate neurons,” *Nature Neuroscience*, vol. 14, no. 12, pp. 1581–1589, 2011.
- [77] N.-Y. So and V. Stuphorn, “Supplementary eye field encodes option and action value for saccades with variable reward,” *Journal of Neurophysiology*, vol. 104, no. 5, pp. 2634–2653, 2010.
- [78] W. Schultz, P. Dayan, and P. R. Montague, “A neural substrate of prediction and reward,” *Science*, vol. 275, no. 5306, pp. 1593–1599, 1997.
- [79] K. Doya, “Complementary roles of basal ganglia and cerebellum in learning and motor control,” *Current Opinion in Neurobiology*, vol. 10, no. 6, pp. 732–739, 2000.
- [80] N. D. Daw and K. Doya, “The computational neurobiology of learning and reward,” *Current Opinion in Neurobiology*, vol. 16, no. 2, pp. 199–204, 2006.
- [81] P. I. Pavlov, “Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex,” *Annals of neurosciences*, vol. 17, no. 3, p. 136, 2010.
- [82] R. R. Bush and F. Mosteller, “A mathematical model for simple learning.,” *Psychological Review*, vol. 58, no. 5, p. 313, 1951.
- [83] R. A. Rescorla, “A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement,” *Current Research and Theory*, pp. 64–99, 1972.
- [84] P. L. Lockwood and M. C. Klein-Flügge, “Computational modelling of social cognition and behaviour—a reinforcement learning primer,” *Social Cognitive and Affective Neuroscience*, vol. 16, no. 8, pp. 761–771, 2021.
- [85] J. Olds and P. Milner, “Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain.,” *Journal of Comparative and Physiological Psychology*, vol. 47, no. 6, p. 419, 1954.

- [86] D. Corbett and R. A. Wise, “Intracranial self-stimulation in relation to the ascending dopaminergic systems of the midbrain: A moveable electrode mapping study,” *Brain Research*, vol. 185, no. 1, pp. 1–15, 1980.
- [87] W. Schultz, “Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey,” *Journal of Neurophysiology*, vol. 56, no. 5, pp. 1439–1461, 1986.
- [88] H. M. Bayer and P. W. Glimcher, “Midbrain dopamine neurons encode a quantitative reward prediction error signal,” *Neuron*, vol. 47, no. 1, pp. 129–141, 2005.
- [89] S. M. McClure, G. S. Berns, and P. R. Montague, “Temporal prediction errors in a passive learning task activate human striatum,” *Neuron*, vol. 38, no. 2, pp. 339–346, 2003.
- [90] N. D. Daw, S. J. Gershman, B. Seymour, P. Dayan, and R. J. Dolan, “Model-based influences on humans’ choices and striatal prediction errors,” *Neuron*, vol. 69, no. 6, pp. 1204–1215, 2011.
- [91] G. R. Yang, M. R. Joglekar, H. F. Song, W. T. Newsome, and X.-J. Wang, “Task representations in neural networks trained to perform many cognitive tasks,” *Nature Neuroscience*, vol. 22, no. 2, pp. 297–306, 2019.
- [92] A. Dezfouli, K. Griffiths, F. Ramos, P. Dayan, and B. W. Balleine, “Models that learn how humans learn: The case of decision-making and its disorders,” *PLoS Computational Biology*, vol. 15, no. 6, e1006903, 2019.
- [93] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [94] A. Dezfouli, H. Ashtiani, O. Ghattas, R. Nock, P. Dayan, and C. S. Ong, “Disentangled behavioural representations.” in *NeurIPS*, 2019, pp. 2251–2260.
- [95] M. Fintz, M. Osadchy, and U. Hertz, “Using deep learning to predict human decisions and using cognitive models to explain deep learning models,” *Scientific Reports*, vol. 12, no. 1, pp. 1–12, 2022.

- [96] A. Bechara, A. R. Damasio, H. Damasio, and S. W. Anderson, “Insensitivity to future consequences following damage to human prefrontal cortex,” *Cognition*, vol. 50, no. 1-3, pp. 7–15, 1994.
- [97] M. Mimura, R. Oeda, and M. Kawamura, “Impaired decision-making in Parkinson’s Disease,” *Parkinsonism & Related Disorders*, vol. 12, no. 3, pp. 169–175, 2006.
- [98] S. W. Kjær, M. B. Callesen, L. Larsen, P. Borghammer, K. Østergaard, and M. F. Damholdt, “Applied strategy in the Iowa Gambling Task: Comparison of individuals with Parkinson’s Disease to healthy controls,” *Journal of Clinical and Experimental Neuropsychology*, vol. 42, no. 5, pp. 425–435, 2020.
- [99] H. Sinz, L. Zamarian, T. Benke, G. Wenning, and M. Delazer, “Impact of ambiguity and risk on decision making in mild Alzheimer’s disease,” *Neuropsychologia*, vol. 46, no. 7, pp. 2043–2055, 2008.
- [100] M. Brière, L. Tocanier, P. Allain, *et al.*, “Decision-making measured by the Iowa Gambling Task in patients with alcohol use disorders choosing harm reduction versus relapse prevention program,” *European Addiction Research*, vol. 25, no. 4, pp. 182–190, 2019.
- [101] F. F. d. Rocha, N. B. Alvarenga, L. Malloy-Diniz, and H. Corrêa, “Decision-making impairment in obsessive-compulsive disorder as measured by the Iowa Gambling Task,” *Arquivos de Neuro-psiquiatria*, vol. 69, no. 4, pp. 642–647, 2011.
- [102] B. Shurman, W. P. Horan, and K. H. Nuechterlein, “Schizophrenia patients demonstrate a distinctive pattern of decision-making impairment on the Iowa Gambling Task,” *Schizophrenia Research*, vol. 72, no. 2-3, pp. 215–224, 2005.
- [103] T. E. Behrens, M. W. Woolrich, M. E. Walton, and M. F. Rushworth, “Learning the value of information in an uncertain world,” *Nature Neuroscience*, vol. 10, no. 9, pp. 1214–1221, 2007.

- [104] M. Browning, T. E. Behrens, G. Jocham, J. X. O'reilly, and S. J. Bishop, "Anxious individuals have difficulty learning the causal statistics of aversive environments," *Nature Neuroscience*, vol. 18, no. 4, pp. 590–596, 2015.
- [105] R. P. Lawson, C. Mathys, and G. Rees, "Adults with autism overestimate the volatility of the sensory environment," *Nature Neuroscience*, vol. 20, no. 9, pp. 1293–1299, 2017.
- [106] C. Manning, J. Kilner, L. Neil, T. Karaminis, and E. Pellicano, "Children on the autism spectrum update their behaviour in response to a volatile environment," *Developmental Science*, vol. 20, no. 5, e12435, 2017.
- [107] E. Pulcu and M. Browning, "Affective bias as a rational response to the statistics of rewards and punishments," *Elife*, vol. 6, e27879, 2017.
- [108] S. Shiffman, "Designing protocols for ecological momentary assessment," *The Science of Real-time Data Capture: Self-reports in Health Research*, pp. 27–53, 2007.
- [109] N. D. Daw *et al.*, "Trial-by-trial data analysis using computational models," *Decision making, Affect, and Learning: Attention and Performance XXIII*, vol. 23, no. 1, 2011.
- [110] R. C. Wilson and A. G. Collins, "Ten simple rules for the computational modeling of behavioral data," *Elife*, vol. 8, e49547, 2019.
- [111] O. Guest and A. E. Martin, "How computational modeling can force theory building in psychological science," *Perspectives on Psychological Science*, vol. 16, no. 4, pp. 789–802, 2021.
- [112] C. Hedge, G. Powell, and P. Sumner, "The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences," *Behavior Research Methods*, vol. 50, no. 3, pp. 1166–1186, 2018.
- [113] N. Haines, P. D. Kvam, L. H. Irving, *et al.*, "Theoretically informed generative models can advance the psychological and brain sciences: Lessons from the reliability paradox," 2020.

- [114] S. Palminteri, V. Wyart, and E. Koechlin, “The importance of falsification in computational cognitive modeling,” *Trends in Cognitive Sciences*, vol. 21, no. 6, pp. 425–433, 2017.
- [115] L. Zhang, L. Lengersdorff, N. Mikus, J. Gläscher, and C. Lamm, “Using reinforcement learning models in social neuroscience: Frameworks, pitfalls and suggestions of best practices,” *Social Cognitive and Affective Neuroscience*, vol. 15, no. 6, pp. 695–707, 2020.
- [116] I. J. Myung, “Tutorial on maximum likelihood estimation,” *Journal of Mathematical Psychology*, vol. 47, no. 1, pp. 90–100, 2003.
- [117] W.-Y. Ahn, A. Krawitz, W. Kim, J. R. Busemeyer, and J. W. Brown, “A model-based fMRI analysis with hierarchical bayesian parameter estimation.,” *Decision*, 2013.
- [118] W.-Y. Ahn, G. Vasilev, S.-H. Lee, *et al.*, “Decision-making in stimulant and opiate addicts in protracted abstinence: Evidence from computational modeling with pure users,” *Frontiers in Psychology*, vol. 5, p. 849, 2014.
- [119] S. Suzuki, R. Adachi, S. Dunne, P. Bossaerts, and J. P. O’Doherty, “Neural mechanisms underlying human consensus decision-making,” *Neuron*, vol. 86, no. 2, pp. 591–602, 2015.
- [120] K. Katahira, “How hierarchical models improve point estimates of model parameters at the individual level,” *Journal of Mathematical Psychology*, vol. 73, pp. 37–58, 2016.
- [121] B. Efron and C. Morris, “Stein’s paradox in statistics,” *Scientific American*, vol. 236, no. 5, pp. 119–127, 1977.
- [122] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian data analysis*. Chapman and Hall/CRC, 1995.
- [123] S. E. Maxwell, H. D. Delaney, and K. Kelley, *Designing experiments and analyzing data: A model comparison perspective*. Routledge, 2017.

- [124] R. McElreath, *Statistical rethinking: A Bayesian course with examples in R and Stan*. Chapman and Hall/CRC, 2020.
- [125] B. Carpenter, A. Gelman, M. D. Hoffman, *et al.*, “Stan: A probabilistic programming language,” *Journal of Statistical Software*, vol. 76, no. 1, 2017.
- [126] D. J. Lunn, A. Thomas, N. Best, and D. Spiegelhalter, “Winbugs—a bayesian modelling framework: Concepts, structure, and extensibility,” *Statistics and Computing*, vol. 10, no. 4, pp. 325–337, 2000.
- [127] M. Plummer, *Jags version 3.3. 0 user manual*, 2012.
- [128] Y. Sakamoto, M. Ishiguro, and G. Kitagawa, “Akaike information criterion statistics,” *Dordrecht, The Netherlands: D. Reidel*, vol. 81, no. 10.5555, p. 26 853, 1986.
- [129] D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. Van Der Linde, “Bayesian measures of model complexity and fit,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 64, no. 4, pp. 583–639, 2002.
- [130] G. Schwarz, “Estimating the dimension of a model,” *The Annals of Statistics*, pp. 461–464, 1978.
- [131] S. Watanabe, “A widely applicable bayesian information criterion,” *ArXiv Preprint*, 2012.
- [132] A. Vehtari, A. Gelman, and J. Gabry, “Practical bayesian model evaluation using leave-one-out cross-validation and waic,” *Statistics and Computing*, vol. 27, no. 5, pp. 1413–1432, 2017.
- [133] S. M. Lynch and B. Western, “Bayesian posterior predictive checks for complex models,” *Sociological Methods & Research*, vol. 32, no. 3, pp. 301–335, 2004.
- [134] R. Levy, R. J. Mislevy, and S. Sinharay, “Posterior predictive model checking for multidimensionality in item response theory,” *Applied Psychological Measurement*, vol. 33, no. 7, pp. 519–537, 2009.

- [135] N. Haines, J. Vassileva, and W.-Y. Ahn, "The outcome-representation learning model: A novel reinforcement learning model of the Iowa Gambling Task," *Cognitive Science*, vol. 42, no. 8, pp. 2534–2561, 2018.
- [136] H. E. Merskey, "Classification of chronic pain: Descriptions of chronic pain syndromes and definitions of pain terms.," *Pain*, 1986.
- [137] S. M. Banks and R. D. Kerns, "Explaining high rates of depression in chronic pain: A diathesis-stress framework.," *Psychological Bulletin*, vol. 119, no. 1, p. 95, 1996.
- [138] R. D. Kerns and M. C. Jacob, "Psychological aspects of back pain," *Baillière's Clinical Rheumatology*, vol. 7, no. 2, pp. 337–356, 1993.
- [139] L. S. Simon, "Relieving pain in america: A blueprint for transforming prevention, care, education, and research," *Journal of Pain & Palliative Care Pharmacotherapy*, vol. 26, no. 2, pp. 197–198, 2012.
- [140] H. Breivik, B. Collett, V. Ventafridda, R. Cohen, and D. Gallacher, "Survey of chronic pain in europe: Prevalence, impact on daily life, and treatment," *European Journal of Pain*, vol. 10, no. 4, pp. 287–333, 2006.
- [141] D. S. Goldberg and S. J. McGee, "Pain as a global public health priority," *BMC Public Health*, vol. 11, no. 1, pp. 1–5, 2011.
- [142] B. E. McGuire, "Chronic pain and cognitive function," *Pain*, vol. 154, no. 7, pp. 964–965, 2013.
- [143] R. Luerding, T. Weigand, U. Bogdahn, and T. Schmidt-Wilcke, "Working memory performance is correlated with local brain morphology in the medial frontal and anterior cingulate cortex in fibromyalgia patients: Structural correlates of pain–cognition interaction," *Brain*, vol. 131, no. 12, pp. 3222–3231, 2008.
- [144] F. Mongini, R. Keller, A. Deregibus, E. Barbalonga, and T. Mongini, "Frontal lobe dysfunction in patients with chronic migraine: A clinical–neuropsychological study," *Psychiatry Research*, vol. 133, no. 1, pp. 101–106, 2005.

- [145] S. Tamburin, A. Maier, S. Schiff, *et al.*, “Cognition and emotional decision-making in chronic low back pain: An ERPs study during Iowa Gambling Task,” *Frontiers in Psychology*, vol. 5, p. 1350, 2014.
- [146] O. Moriarty, N. Ruane, D. O’Gorman, *et al.*, “Cognitive impairment in patients with chronic neuropathic or radicular pain: An interaction of pain and age,” *Frontiers in Behavioral Neuroscience*, vol. 11, p. 100, 2017.
- [147] A. V. Apkarian, Y. Sosa, S. Sonty, *et al.*, “Chronic back pain is associated with decreased prefrontal and thalamic gray matter density,” *Journal of Neuroscience*, vol. 24, no. 46, pp. 10 410–10 415, 2004.
- [148] F. Cauda, S. Palermo, T. Costa, *et al.*, “Gray matter alterations in chronic pain: A network-oriented meta-analytic approach,” *NeuroImage: Clinical*, vol. 4, pp. 676–686, 2014.
- [149] C. Yuan, H. Shi, P. Pan, *et al.*, “Gray matter abnormalities associated with chronic back pain,” *The Clinical Journal of Pain*, vol. 33, no. 11, pp. 983–990, 2017.
- [150] J. M. Glass, D. A. Williams, M.-L. Fernandez-Sanchez, *et al.*, “Executive function in chronic pain patients and healthy controls: Different cortical activation during response inhibition in fibromyalgia,” *The journal of pain*, vol. 12, no. 12, pp. 1219–1229, 2011.
- [151] A. Bechara and A. R. Damasio, “The somatic marker hypothesis: A neural theory of economic decision,” *Games and Economic Behavior*, vol. 52, no. 2, pp. 336–372, 2005.
- [152] M. I. Posner, M. K. Rothbart, B. E. Sheese, and Y. Tang, “The anterior cingulate gyrus and the mechanism of self-regulation,” *Cognitive, Affective, & Behavioral Neuroscience*, vol. 7, no. 4, pp. 391–395, 2007.
- [153] M. F. Rushworth, M. P. Noonan, E. D. Boorman, M. E. Walton, and T. E. Behrens, “Frontal cortex and reward-guided learning and decision-making,” *Neuron*, vol. 70, no. 6, pp. 1054–1069, 2011.

- [154] P. W. Frankland, B. Bontempi, L. E. Talton, L. Kaczmarek, and A. J. Silva, “The involvement of the anterior cingulate cortex in remote contextual fear memory,” *Science*, vol. 304, no. 5672, pp. 881–883, 2004.
- [155] B. Naylor, N. Hesam-Shariati, J. H. McAuley, *et al.*, “Reduced glutamate in the medial prefrontal cortex is associated with emotional and cognitive dysregulation in people with chronic pain,” *Frontiers in Neurology*, vol. 10, p. 1110, 2019.
- [156] A. Bechara, H. Damasio, A. R. Damasio, and G. P. Lee, “Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making,” *Journal of Neuroscience*, vol. 19, no. 13, pp. 5473–5481, 1999.
- [157] A. Bechara, D. Tranel, and H. Damasio, “Characterization of the decision-making deficit of patients with ventromedial prefrontal cortex lesions,” *Brain*, vol. 123, no. 11, pp. 2189–2202, 2000.
- [158] P. Cavedini, G. Riboldi, A. D’Annunzi, P. Belotti, M. Cisima, and L. Bellodi, “Decision-making heterogeneity in obsessive-compulsive disorder: Ventromedial prefrontal cortex function predicts different treatment outcomes,” *Neuropsychologia*, vol. 40, no. 2, pp. 205–211, 2002.
- [159] A. Bechara and H. Damasio, “Decision-making and addiction (part i): Impaired activation of somatic states in substance dependent individuals when pondering decisions with negative future consequences,” *Neuropsychologia*, vol. 40, no. 10, pp. 1675–1689, 2002.
- [160] A. Verdejo-García, F. Lopez-Torrecillas, E. P. Calandre, A. Delgado-Rodríguez, and A. Bechara, “Executive function and decision-making in women with fibromyalgia,” *Archives of Clinical Neuropsychology*, vol. 24, no. 1, pp. 113–122, 2009.
- [161] L. E. Hess, A. Haimovici, M. A. Muñoz, and P. Montoya, “Beyond pain: Modeling decision-making deficits in chronic pain,” *Frontiers in Behavioral Neuroscience*, vol. 8, p. 263, 2014.

- [162] E. Yechiam, J. E. Kanz, A. Bechara, *et al.*, “Neurocognitive deficits related to poor decision making in people behind bars,” *Psychonomic Bulletin & Review*, vol. 15, no. 1, pp. 44–51, 2008.
- [163] T. Mendoza, T. Mayne, D. Rublee, and C. Cleeland, “Reliability and validity of a modified brief pain inventory short form in patients with osteoarthritis,” *European Journal of Pain*, vol. 10, no. 4, pp. 353–361, 2006.
- [164] W.-Y. Ahn, J. R. Busemeyer, E.-J. Wagenmakers, and J. C. Stout, “Comparison of decision learning models using the generalization criterion method,” *Cognitive Science*, vol. 32, no. 8, pp. 1376–1402, 2008.
- [165] H. Steingroever, T. Pachur, M. Šm’ira, and M. D. Lee, “Bayesian techniques for analyzing group differences in the Iowa Gambling Task: A case study of intuitive and deliberate decision-makers,” *Psychonomic Bulletin & Review*, vol. 25, no. 3, pp. 951–970, 2018.
- [166] J. R. Busemeyer and J. C. Stout, “A contribution of cognitive decision models to clinical assessment: Decomposing performance on the bechara gambling task.” *Psychological Assessment*, vol. 14, no. 3, p. 253, 2002.
- [167] W.-Y. Ahn, N. Haines, and L. Zhang, “Revealing neurocomputational mechanisms of reinforcement learning and decision-making with the hbayesdm package,” *Computational Psychiatry (Cambridge, Mass.)*, vol. 1, p. 24, 2017.
- [168] A. V. Apkarian, Y. Sosa, B. R. Krauss, *et al.*, “Chronic pain patients are impaired on an emotional decision-making task,” *Pain*, vol. 108, no. 1-2, pp. 129–136, 2004.
- [169] B. D. Dick and S. Rashiq, “Disruption of attention and working memory traces in individuals with chronic pain,” *Anesthesia & Analgesia*, vol. 104, no. 5, pp. 1223–1229, 2007.
- [170] L. L. Jorge, C. Gerard, and M. Revel, “Evidences of memory dysfunction and maladaptive coping in chronic low back pain and rheumatoid arthritis patients: Challenges for rehabilitation,” *Eur J Phys Rehabil Med*, vol. 45, no. 4, pp. 469–477, 2009.

- [171] X. Liu, L. Li, F. Tang, S. Wu, and Y. Hu, “Memory impairment in chronic pain patients and the related neuropsychological mechanisms: A review,” *Acta Neuropsychiatrica*, vol. 26, no. 4, pp. 195–201, 2014.
- [172] G. DeMaagd and A. Philip, “Parkinson’s Disease and its management: Part 1: Disease entity, risk factors, pathophysiology, clinical presentation, and diagnosis,” *Pharmacy and Therapeutics*, vol. 40, no. 8, p. 504, 2015.
- [173] L. M. De Lau and M. M. Breteler, “Epidemiology of Parkinson’s Disease,” *The Lancet Neurology*, vol. 5, no. 6, pp. 525–535, 2006.
- [174] G. Morris, A. Nevet, D. Arkadir, E. Vaadia, and H. Bergman, “Midbrain dopamine neurons encode decisions for future action,” *Nature Neuroscience*, vol. 9, no. 8, pp. 1057–1063, 2006.
- [175] A. Lak, W. R. Stauffer, and W. Schultz, “Dopamine neurons learn relative chosen value from probabilistic rewards,” *Elife*, vol. 5, e18044, 2016.
- [176] D. Shohamy, C. Myers, S. Onlaor, and M. Gluck, “Role of the basal ganglia in category learning: How do patients with Parkinson’s Disease learn?” *Behavioral Neuroscience*, vol. 118, no. 4, p. 676, 2004.
- [177] M. J. Frank, L. C. Seeberger, and R. C. O’reilly, “By carrot or by stick: Cognitive reinforcement learning in parkinsonism,” *Science*, vol. 306, no. 5703, pp. 1940–1943, 2004.
- [178] D. A. Peterson, C. Elliott, D. D. Song, S. Makeig, T. J. Sejnowski, and H. Poizner, “Probabilistic reversal learning is impaired in Parkinson’s Disease,” *Neuroscience*, vol. 163, no. 4, pp. 1092–1101, 2009.
- [179] D. Shohamy, C. E. Myers, K. D. Geghman, J. Sage, and M. A. Gluck, “L-dopa impairs learning, but spares generalization, in Parkinson’s Disease,” *Neuropsychologia*, vol. 44, no. 5, pp. 774–784, 2006.
- [180] M. Jahanshahi, L. Wilkinson, H. Gahir, A. Dharminda, and D. A. Lagnado, “Medication impairs probabilistic classification learning in Parkinson’s Disease,” *Neuropsychologia*, vol. 48, no. 4, pp. 1096–1103, 2010.

- [181] A. Lees, “The on-off phenomenon,” *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 52, no. Suppl, pp. 29–37, 1989.
- [182] R. Cools, L. Altamirano, and M. D’Esposito, “Reversal learning in Parkinson’s Disease depends on medication status and outcome valence,” *Neuropsychologia*, vol. 44, no. 10, pp. 1663–1673, 2006.
- [183] R. Swinson, R. Rogers, B. Sahakian, B. Summers, C. Polkey, and T. Robbins, “Probabilistic learning and reversal deficits in patients with Parkinson’s Disease or frontal or temporal lobe lesions: Possible adverse effects of dopaminergic medication,” *Neuropsychologia*, vol. 38, no. 5, pp. 596–612, 2000.
- [184] M. J. Frank, “Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and nonmedicated parkinsonism,” *Journal of Cognitive Neuroscience*, vol. 17, no. 1, pp. 51–72, 2005.
- [185] S. Graef, G. Biele, L. K. Krugel, *et al.*, “Differential influence of levodopa on reward-based learning in Parkinson’s Disease,” *Frontiers in Human Neuroscience*, vol. 4, p. 169, 2010.
- [186] P. Smittenaar, H. Chase, E. Aarts, B. Nusslein, B. Bloem, and R. Cools, “Decomposing effects of dopaminergic medication in Parkinson’s Disease on probabilistic action selection–learning or performance?” *European Journal of Neuroscience*, vol. 35, no. 7, pp. 1144–1151, 2012.
- [187] J. P. Grogan, D. Tsivos, L. Smith, *et al.*, “Effects of dopamine on reinforcement learning and consolidation in Parkinson’s Disease,” *Elife*, vol. 6, e26801, 2017.
- [188] M. H. Timmer, G. Sescousse, M. E. van der Schaaf, R. A. Esselink, and R. Cools, “Reward learning deficits in Parkinson’s Disease depend on depression,” *Psychological Medicine*, vol. 47, no. 13, pp. 2302–2311, 2017.
- [189] J. Aylward, V. Valton, W.-Y. Ahn, *et al.*, “Altered learning under uncertainty in unmedicated mood and anxiety disorders,” *Nature Human Behaviour*, vol. 3, no. 10, pp. 1116–1123, 2019.

- [190] V. M. Brown, L. Zhu, A. Solway, *et al.*, “Reinforcement learning disruptions in individuals with depression and sensitivity to symptom change following cognitive behavioral therapy,” *JAMA Psychiatry*, vol. 78, no. 10, pp. 1113–1122, 2021.
- [191] A. C. Pike and O. J. Robinson, “Reinforcement learning in patients with mood and anxiety disorders vs control individuals: A systematic review and meta-analysis,” *JAMA Psychiatry*, 2022.
- [192] D. Mukherjee, A. L. Filipowicz, K. Vo, T. D. Satterthwaite, and J. W. Kable, “Reward and punishment reversal-learning in major depressive disorder,” *Journal of Abnormal Psychology*, vol. 129, no. 8, p. 810, 2020.
- [193] K. R. Chaudhuri, D. G. Healy, and A. H. Schapira, “Non-motor symptoms of Parkinson’s Disease: Diagnosis and management,” *The Lancet Neurology*, vol. 5, no. 3, pp. 235–245, 2006.
- [194] A. H. Schapira, K. Chaudhuri, and P. Jenner, “Non-motor features of Parkinson Disease,” *Nature Reviews Neuroscience*, vol. 18, no. 7, pp. 435–450, 2017.
- [195] E. M. Khedr, A. A. Abdelrahman, Y. Elserogy, A. F. Zaki, and A. Gamea, “Depression and anxiety among patients with Parkinson’s Disease: Frequency, risk factors, and impact on quality of life,” *The Egyptian Journal of Neurology, Psychiatry and Neurosurgery*, vol. 56, no. 1, pp. 1–9, 2020.
- [196] D. Aarsland, L. Marsh, and A. Schrag, “Neuropsychiatric symptoms in Parkinson’s Disease,” *Movement Disorders: Official Journal of the Movement Disorder Society*, vol. 24, no. 15, pp. 2175–2186, 2009.
- [197] H. S. Mayberg and D. H. Solomon, “Depression in Parkinson’s Disease: A biochemical and organic viewpoint,” *Advances in Neurology*, vol. 65, pp. 49–60, 1995.
- [198] C. Gagne, O. Zika, P. Dayan, and S. J. Bishop, “Impaired adaptation of learning to contingency volatility in internalizing psychopathology,” *Elife*, vol. 9, e61387, 2020.

- [199] I. Litvan, K. P. Bhatia, D. J. Burn, *et al.*, “Movement disorders society scientific issues committee report: Sic task force appraisal of clinical diagnostic criteria for parkinsonian disorders.,” *Movement Disorders: Official Journal of the Movement Disorder Society*, vol. 18, no. 5, pp. 467–486, 2003.
- [200] M. J. Müller and A. Dragicevic, “Standardized rater training for the hamilton depression rating scale (hamd-17) in psychiatric novices,” *Journal of Affective Disorders*, vol. 77, no. 1, pp. 65–69, 2003.
- [201] V. Peto, C. Jenkinson, R. Fitzpatrick, and R. Greenhall, “The development and validation of a short measure of functioning and well being for individuals with Parkinson’s Disease,” *Quality of Life Research*, vol. 4, no. 3, pp. 241–248, 1995.
- [202] P. Naarding, A. F. Leentjens, F. van Kooten, and F. R. Verhey, “Disease-specific properties of the hamilton rating scale for depression in patients with stroke, Alzheimer’s dementia, and Parkinson’s Disease,” *The Journal of Neuropsychiatry and Clinical Neurosciences*, vol. 14, no. 3, pp. 329–334, 2002.
- [203] A. F. Leentjens, F. R. Verhey, R. Lousberg, H. Spitsbergen, and F. W. Wilmkink, “The validity of the hamilton and montgomery-Åsberg depression rating scales as screening and diagnostic tools for depression in Parkinson’s Disease,” *International Journal of Geriatric Psychiatry*, vol. 15, no. 7, pp. 644–649, 2000.
- [204] N. Dragašević, A. Potrebić, A. Damjanović, E. Stefanova, and V. S. Kostić, “Therapeutic efficacy of bilateral prefrontal slow repetitive transcranial magnetic stimulation in depressed patients with Parkinson’s Disease: An open study,” *Movement Disorders: Official Journal of the Movement Disorder Society*, vol. 17, no. 3, pp. 528–532, 2002.
- [205] F. Fregni, C. Santos, M. Myczkowski, *et al.*, “Repetitive transcranial magnetic stimulation is as effective as fluoxetine in the treatment of depression in patients with Parkinson’s Disease,” *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 75, no. 8, pp. 1171–1174, 2004.

- [206] A. Kummer, F. Cardoso, and A. L. Teixeira, “Generalized anxiety disorder and the hamilton anxiety rating scale in Parkinson’s Disease,” *Arquivos de Neuro-psiquiatria*, vol. 68, pp. 495–501, 2010.
- [207] K. R. Chaudhuri, A. Rizos, C. Trenkwalder, *et al.*, “King’s Parkinson’s Disease pain scale, the first scale for pain in pd: An international validation,” *Movement Disorders*, vol. 30, no. 12, pp. 1623–1631, 2015.
- [208] C. H. Donahue and D. Lee, “Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex,” *Nature Neuroscience*, vol. 18, no. 2, pp. 295–301, 2015.
- [209] A. Gelman and D. B. Rubin, “Inference from iterative simulation using multiple sequences,” *Statistical Science*, pp. 457–472, 1992.
- [210] G. S. Berns and E. Bell, “Striatal topography of probability and magnitude information for decisions under uncertainty,” *NeuroImage*, vol. 59, no. 4, pp. 3166–3172, 2012.
- [211] E. Pulcu and M. Browning, “The misestimation of uncertainty in affective disorders,” *Trends in Cognitive Sciences*, vol. 23, no. 10, pp. 865–875, 2019.
- [212] M. Cohen and C. Ranganath, “Behavioral and neural predictors of upcoming decisions,” *Cognitive, Affective, & Behavioral Neuroscience*, vol. 5, no. 2, pp. 117–126, 2005.
- [213] D. Baguley, D. McFerran, and D. Hall, “Tinnitus,” *The Lancet*, vol. 382, no. 9904, pp. 1600–1607, 2013.
- [214] J. Shargorodsky, G. C. Curhan, and W. R. Farwell, “Prevalence and characteristics of tinnitus among us adults,” *The American Journal of Medicine*, vol. 123, no. 8, pp. 711–718, 2010.
- [215] N. A. Stohler, D. Reinau, S. S. Jick, D. Bodmer, and C. R. Meier, “A study on the epidemiology of tinnitus in the united kingdom,” *Clinical Epidemiology*, vol. 11, p. 855, 2019.

- [216] W. Schlee, S. Hølleland, J. Bulla, *et al.*, “The effect of environmental stressors on tinnitus: A prospective longitudinal study on the impact of the covid-19 pandemic,” *Journal of Clinical Medicine*, vol. 9, no. 9, p. 2756, 2020.
- [217] B. Langguth, “A review of tinnitus symptoms beyond ’ringing in the ears’: A call to action,” *Current Medical Research and Opinion*, vol. 27, no. 8, pp. 1635–1643, 2011.
- [218] A. B. Elgoyhen, B. Langguth, D. De Ridder, and S. Vanneste, “Tinnitus: Perspectives from human neuroimaging,” *Nature Reviews Neuroscience*, vol. 16, no. 10, pp. 632–642, 2015.
- [219] D. De Ridder, W. Schlee, S. Vanneste, *et al.*, “Tinnitus and tinnitus disorder: Theoretical and operational definitions (an international multidisciplinary proposal),” *Progress in Brain Research*, vol. 260, pp. 1–25, 2021.
- [220] C. Lanting, E. De Kleine, and P. Van Dijk, “Neural activity underlying tinnitus generation: Results from PET and fMRI,” *Hearing Research*, vol. 255, no. 1-2, pp. 1–13, 2009.
- [221] A. Crippa, C. P. Lanting, P. Van Dijk, and J. B. Roerdink, “A diffusion tensor imaging study on the auditory system and tinnitus,” *The Open Neuroimaging Journal*, vol. 4, p. 16, 2010.
- [222] S. Rossiter, C. Stevens, and G. Walker, “Tinnitus and its effect on working memory and attention,” 2006.
- [223] K. J. Trevis, N. M. McLachlan, and S. J. Wilson, “Cognitive mechanisms in chronic tinnitus: Psychological markers of a failure to switch attention,” *Frontiers in Psychology*, vol. 7, p. 1262, 2016.
- [224] K. J. Pierce, D. Kallogjeri, J. F. Piccirillo, K. S. Garcia, J. E. Nicklaus, and H. Burton, “Effects of severe bothersome tinnitus on cognitive function measured with standardized tests,” *Journal of Clinical and Experimental Neuropsychology*, vol. 34, no. 2, pp. 126–134, 2012.

- [225] S. K. Das, A. Wineland, D. Kallogjeri, and J. F. Piccirillo, "Cognitive speed as an objective measure of tinnitus," *The Laryngoscope*, vol. 122, no. 11, pp. 2533–2538, 2012.
- [226] S. Vanneste, M. Plazier, E. Van Der Loo, P. Van de Heyning, M. Congedo, and D. De Ridder, "The neural correlates of tinnitus-related distress," *Neuroimage*, vol. 52, no. 2, pp. 470–480, 2010.
- [227] P. Domenech and E. Koechlin, "Executive control and decision-making in the prefrontal cortex," *Current Opinion in Behavioral Sciences*, vol. 1, pp. 101–106, 2015.
- [228] R. Hallam, L. McKenna, and L. Shurlock, "Tinnitus impairs cognitive efficiency," *International Journal of Audiology*, vol. 43, no. 4, pp. 218–226, 2004.
- [229] G. Andersson, J. Eriksson, L.-G. Lundh, and L. Lyttkens, "Tinnitus and cognitive interference: A stroop paradigm study," *Journal of Speech, Language, and Hearing Research*, vol. 43, no. 5, pp. 1168–1173, 2000.
- [230] G. Andersson, R. Bakhsh, L. Johansson, V. Kaldo, and P. Carlbring, "Stroop facilitation in tinnitus patients: An experiment conducted via the world wide web," *CyberPsychology & Behavior*, vol. 8, no. 1, pp. 32–38, 2005.
- [231] K.-M. Holgers, S. Zöger, and K. Svedlund, "Predictive factors for development of severe tinnitus suffering-further characterisation: Factores predictivos para el desarrollo de tinitus severo que sufren una caracterización adicional," *International Journal of Audiology*, vol. 44, no. 10, pp. 584–592, 2005.
- [232] A. R. Fetoni, D. Lucidi, E. De Corso, A. Fiorita, G. Conti, and G. Paludetti, "Relationship between subjective tinnitus perception and psychiatric discomfort," *The International Tinnitus Journal*, vol. 20, no. 2, pp. 76–82, 2016.
- [233] B. Langguth, M. Landgrebe, T. Kleinjung, G. P. Sand, and G. Hajak, "Tinnitus and depression," *The World Journal of Biological Psychiatry*, vol. 12, no. 7, pp. 489–500, 2011.

- [234] R. Pryss, W. Schlee, B. Langguth, and M. Reichert, “Mobile crowdsensing services for tinnitus assessment and patient feedback,” in *2017 IEEE International Conference on AI & Mobile Services (AIMS)*, IEEE, 2017, pp. 22–29.
- [235] R. L. Goldberg, M. L. Piccirillo, J. Nicklaus, *et al.*, “Evaluation of ecological momentary assessment for tinnitus severity,” *JAMA Otolaryngology–Head & Neck Surgery*, vol. 143, no. 7, pp. 700–706, 2017.
- [236] M. P. Lourenco, J. Simoes, J. W. Vlaeyen, and R. F. Cima, “The daily experience of subjective tinnitus: Ecological momentary assessment versus end-of-day diary,” *Ear and Hearing*, vol. 43, no. 1, p. 45, 2022.
- [237] M. B. Wilson, D. Kallogjeri, C. N. Joplin, *et al.*, “Ecological momentary assessment of tinnitus using smartphone technology: A pilot study,” *Otolaryngology–Head and Neck Surgery*, vol. 152, no. 5, pp. 897–903, 2015.
- [238] W. Schlee, R. C. Pryss, T. Probst, *et al.*, “Measuring the moment-to-moment variability of tinnitus: The trackyourtinnitus smart phone app,” *Frontiers in Aging Neuroscience*, vol. 8, p. 294, 2016.
- [239] M. McCarthy, L. Zhang, G. Monacelli, T. Ward, *et al.*, “Using methods from computational decision-making to predict nonadherence to fitness goals: Protocol for an observational study,” *JMIR Research Protocols*, vol. 10, no. 11, e29758, 2021.
- [240] E. Genitsaridi, M. Partyka, S. Gallus, *et al.*, “Standardised profiling for tinnitus research: The european school for interdisciplinary tinnitus research screening questionnaire (esit-sq),” *Hearing Research*, vol. 377, pp. 353–359, 2019.
- [241] C. D. Spielberger, F. Gonzalez-Reigosa, A. Martinez-Urrutia, L. F. Natalicio, and D. S. Natalicio, “The state-trait anxiety inventory,” *Revista Interamericana de Psicologia/Interamerican Journal of Psychology*, vol. 5, no. 3 & 4, 1971.
- [242] P. Bech, N.-A. Rasmussen, L. R. Olsen, V. Noerholm, and W. Abildgaard, “The sensitivity and specificity of the major depression inventory, using the present state

- examination as the index of diagnostic validity,” *Journal of Affective Disorders*, vol. 66, no. 2-3, pp. 159–164, 2001.
- [243] W. Hiller and G. Goebel, “Rapid assessment of tinnitus-related psychological distress using the mini-tq,” *Int J Audiol*, vol. 43, no. 10, pp. 600–604, 2004.
- [244] A. Bunevicius, M. Staniute, J. Brozaitiene, V. J. Pop, J. Neverauskas, and R. Bunevicius, “Screening for anxiety disorders in patients with coronary artery disease,” *Health and Quality of Life Outcomes*, vol. 11, no. 1, pp. 1–9, 2013.
- [245] P. Cuijpers, J. Dekker, A. Noteboom, N. Smits, and J. Peen, “Sensitivity and specificity of the major depression inventory in outpatients,” *BMC Psychiatry*, vol. 7, no. 1, pp. 1–6, 2007.
- [246] A. B. Konova, S. Lopez-Guzman, A. Urmanche, *et al.*, “Computational markers of risky decision-making for identification of temporal windows of vulnerability to opioid use in a real-world clinical setting,” *JAMA Psychiatry*, vol. 77, no. 4, pp. 368–377, 2020.
- [247] E. E. Alvarez, S. Hafezi, D. Bonagura, E. M. Kleiman, and A. B. Konova, “A proof-of-concept ecological momentary assessment study of day-level dynamics in value-based decision-making in opioid addiction,” *Frontiers in Psychiatry*, vol. 13, 2022.
- [248] J. M. Bhatt, N. Bhattacharyya, and H. W. Lin, “Relationships between tinnitus and the prevalence of anxiety and depression,” *The Laryngoscope*, vol. 127, no. 2, pp. 466–469, 2017.
- [249] C. Tikkinen-Piri, A. Rohunen, and J. Markkula, “Eu general data protection regulation: Changes and implications for personal data collecting companies,” *Computer Law & Security Review*, vol. 34, no. 1, pp. 134–153, 2018.
- [250] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik, “Federated optimization: Distributed machine learning for on-device intelligence,” *ArXiv Preprint*, 2016.
- [251] A. Hard, K. Rao, R. Mathews, *et al.*, “Federated learning for mobile keyboard prediction,” *ArXiv Preprint*, 2018.

- [252] M. J. Sheller, G. A. Reina, B. Edwards, J. Martin, and S. Bakas, “Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation,” in *International MICCAI Brainlesion Workshop*, Springer, 2018, pp. 92–104.
- [253] K. Chang, N. Balachandar, C. Lam, *et al.*, “Distributed deep learning networks among institutions for medical imaging,” *Journal of the American Medical Informatics Association*, vol. 25, no. 8, pp. 945–954, 2018.
- [254] M. J. Sheller, B. Edwards, G. A. Reina, *et al.*, “Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data,” *Scientific Reports*, vol. 10, no. 1, pp. 1–12, 2020.
- [255] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [256] H. Steingroever, D. J. Fridberg, A. Horstmann, *et al.*, “Data from 617 healthy participants performing the Iowa Gambling Task: A “many labs” collaboration,” *Journal of Open Psychology Data*, vol. 3, no. 1, pp. 340–353, 2015.
- [257] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, “Federated learning with non-iid data,” *ArXiv Preprint*, 2018.
- [258] F. A. Gers, J. Schmidhuber, and F. Cummins, “Learning to forget: Continual prediction with LSTM,” *Neural Computation*, vol. 12, no. 10, pp. 2451–2471, 2000.
- [259] K. Cho, B. Van Merriënboer, C. Gulcehre, *et al.*, “Learning phrase representations using rnn encoder-decoder for statistical machine translation,” *ArXiv Preprint*, 2014.
- [260] M. Abadi, A. Agarwal, P. Barham, *et al.*, “Tensorflow: Large-scale machine learning on heterogeneous distributed systems,” *ArXiv Preprint*, 2016.
- [261] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *ArXiv Preprint*, 2014.

- [262] G. H. Lee and S.-Y. Shin, “Federated learning on clinical benchmark data: Performance assessment,” *Journal of Medical Internet Research*, vol. 22, no. 10, e20891, 2020.
- [263] X. Li, Y. Gu, N. Dvornek, L. H. Staib, P. Ventola, and J. S. Duncan, “Multi-site fMRI analysis using privacy-preserving federated learning and domain adaptation: Abide results,” *Medical Image Analysis*, vol. 65, p. 101 765, 2020.
- [264] J. C. Liu, J. Goetz, S. Sen, and A. Tewari, “Learning from others without sacrificing privacy: Simulation comparing centralized and federated machine learning on mobile health data,” *JMIR MHealth and uHealth*, vol. 9, no. 3, e23728, 2021.
- [265] F. K. Dankar, N. Madathil, S. K. Dankar, and S. Boughorbel, “Privacy-preserving analysis of distributed biomedical data: Designing efficient and secure multiparty computations using distributed statistical learning theory,” *JMIR Medical Informatics*, vol. 7, no. 2, e12702, 2019.
- [266] M. S. Ozdayi, M. Kantarcioglu, and R. Iyer, “Improving accuracy of federated learning in non-iid settings,” *ArXiv Preprint*, 2020.
- [267] H. Wang, Z. Kaplan, D. Niu, and B. Li, “Optimizing federated learning on non-iid data with reinforcement learning,” in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, IEEE, 2020, pp. 1698–1707.
- [268] D. Lee, “Decision making: From neuroscience to psychiatry,” *Neuron*, vol. 78, no. 2, pp. 233–248, 2013.
- [269] C. M. Gillan and N. D. Daw, “Taking psychiatry research online,” *Neuron*, vol. 91, no. 1, pp. 19–23, 2016.
- [270] T. Wise and R. J. Dolan, “Associations between aversive learning processes and transdiagnostic psychiatric symptoms in a general population sample,” *Nature Communications*, vol. 11, no. 1, pp. 1–13, 2020.
- [271] C. M. Gillan, M. Kosinski, R. Whelan, E. A. Phelps, and N. D. Daw, “Characterizing a psychiatric symptom dimension related to deficits in goal-directed control,” *Elife*, vol. 5, e11305, 2016.

- [272] M. Rouault, T. Seow, C. M. Gillan, and S. M. Fleming, “Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance,” *Biological Psychiatry*, vol. 84, no. 6, pp. 443–451, 2018.
- [273] S. Triantafyllou, S. Saeb, E. G. Lattie, D. C. Mohr, K. P. Kording, *et al.*, “Relationship between sleep quality and mood: Ecological momentary assessment study,” *JMIR Mental Health*, vol. 6, no. 3, e12613, 2019.
- [274] E. R. Weitzman, “Poor mental health, depression, and associations with alcohol consumption, harm, and abuse in a national sample of young adults in college,” *The Journal of Nervous and Mental Disease*, vol. 192, no. 4, pp. 269–277, 2004.
- [275] E. Jane-Llopis and I. Matytsina, “Mental health and alcohol, drugs and tobacco: A review of the comorbidity between mental disorders and the use of alcohol, tobacco and illicit drugs,” *Drug and Alcohol Review*, vol. 25, no. 6, pp. 515–536, 2006.
- [276] A. S. S. de Siqueira, M. K. Flaks, M. M. Biella, S. Mauer, M. K. Borges, and I. Aprahamian, “Decision making assessed by the Iowa Gambling Task and major depressive disorder a systematic review,” *Dementia & Neuropsychologia*, vol. 12, pp. 250–255, 2018.
- [277] T. Soshi, M. Nagamine, E. Fukuda, and A. Takeuchi, “Pre-specified anxiety predicts future decision-making performances under different temporally constrained conditions,” *Frontiers in Psychology*, vol. 10, p. 1544, 2019.
- [278] G. Monacelli, L. Zhang, W. Schlee, B. Langguth, T. E. Ward, and T. B. Murphy, “Adaptive data collection for intra-individual studies affected by adherence,” *ArXiv Preprint*, 2022.
- [279] T. R. Insel, “Digital phenotyping: A global tool for psychiatry,” *World Psychiatry*, vol. 17, no. 3, p. 276, 2018.
- [280] D. Groh, “The temptation of conspiracy theory, or: Why do bad things happen to good people? part ii: Case studies,” in *Changing Conceptions of Conspiracy*, Springer, 1987, pp. 15–37.

- [281] G. Sedek, "Conspiracy stereotypes of jews during systemic transformation in poland," *International Journal of Sociology*, vol. 35, no. 1, pp. 40–64, 2005.
- [282] M. Zonis and C. M. Joseph, "Conspiracy thinking in the middle east," *Political Psychology*, pp. 443–459, 1994.
- [283] R. Brotherton and S. Eser, "Bored to fears: Boredom proneness, paranoia, and conspiracy theories," *Personality and Individual Differences*, vol. 80, pp. 1–5, 2015.
- [284] P. J. Leman and M. Cinnirella, "Beliefs in conspiracy theories and the need for cognitive closure," *Frontiers in psychology*, vol. 4, p. 378, 2013.
- [285] J.-W. van Prooijen and M. Acker, "The influence of control on belief in conspiracy theories: Conceptual and applied extensions," *Applied Cognitive Psychology*, vol. 29, no. 5, pp. 753–761, 2015.
- [286] V. Swami, R. Coles, S. Stieger, *et al.*, "Conspiracist ideation in britain and austria: Evidence of a monological belief system and associations between individual psychological differences and real-world and fictitious conspiracy theories," *British Journal of Psychology*, vol. 102, no. 3, pp. 443–463, 2011.
- [287] K. M. Douglas and R. M. Sutton, "Does it take one to know one? endorsement of conspiracy theories is influenced by personal willingness to conspire," *British Journal of Social Psychology*, vol. 50, no. 3, pp. 544–552, 2011.
- [288] M. Grzesiak-Feldman, "Conspiracy thinking and state-trait anxiety in young polish adults," *Psychological Reports*, vol. 100, no. 1, pp. 199–202, 2007.
- [289] M. Grzesiak-Feldman, "The effect of high-anxiety situations on conspiracy thinking," *Current Psychology*, vol. 32, no. 1, pp. 100–118, 2013.
- [290] M. Grzesiak-Feldman and M. Irzycka, "Right-wing authoritarianism and conspiracy thinking in a polish sample," *Psychological reports*, vol. 105, no. 2, pp. 389–393, 2009.

- [291] J.-W. Van Prooijen, K. M. Douglas, and C. De Inocencio, “Connecting the dots: Illusory pattern perception predicts belief in conspiracies and the supernatural,” *European journal of social psychology*, vol. 48, no. 3, pp. 320–335, 2018.