Society for NeuroEconomics

NEUROSCIENCE • PSYCHOLOGY • ECONOMICS

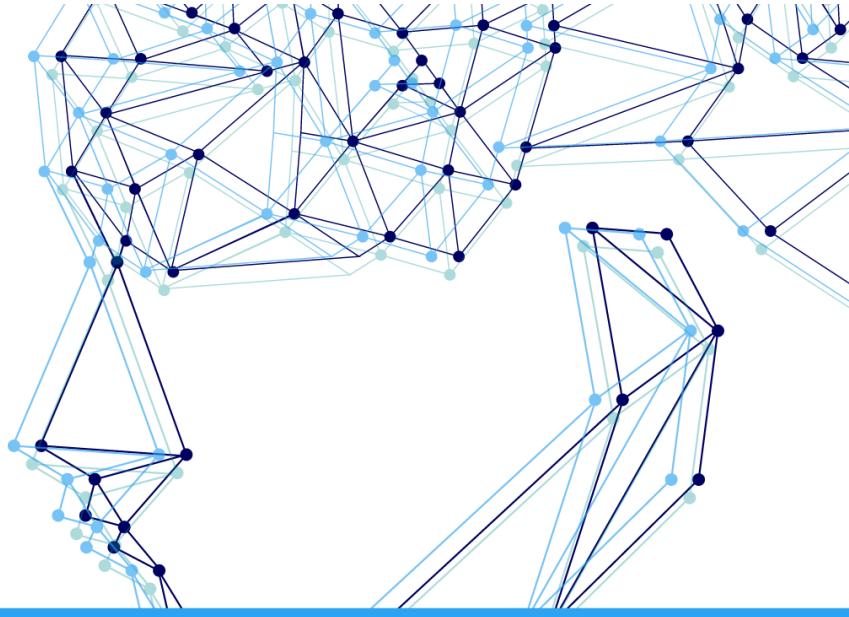

ABSTRACT BOOK

19th Annual Meeting

September 29 – October 1, 2021

#vSNE2021

Society for NeuroEconomics
Program at a Glance
2021 vSNE Conference

Los Angeles	New York	London	Paris	Tokyo	Sydney
PDT	EDT	BST	CEST	JST	AEST
5:00	8:00	13:00	14:00	21:00	22:00
5:30	8:30	13:30	14:30	21:30	22:30
6:00	9:00	14:00	15:00	22:00	23:00
6:30	9:30	14:30	15:30	22:30	23:30
7:00	10:00	15:00	16:00	23:00	0:00
7:30	10:30	15:30	16:30	23:30	0:30
8:00	11:00	16:00	17:00	0:00	1:00
8:30	11:30	16:30	17:30	0:30	1:30
9:00	12:00	17:00	18:00	1:00	2:00
9:30	12:30	17:30	18:30	1:30	2:30
10:00	13:00	18:00	19:00	2:00	3:00
10:30	13:30	18:30	19:30	2:30	3:30
11:00	14:00	19:00	20:00	3:00	4:00
11:30	14:30	19:30	20:30	3:30	4:30
12:00	15:00	20:00	21:00	4:00	5:00
12:30	15:30	20:30	21:30	4:30	5:30
13:00	16:00	21:00	22:00	5:00	6:00
13:30	16:30	21:30	22:30	5:30	6:30
14:00	17:00	22:00	23:00	6:00	7:00
14:30	17:30	22:30	23:30	6:30	7:30
15:00	18:00	23:00	0:00	7:00	8:00
15:30	18:30	23:30	0:30	7:30	8:30
16:00	19:00	0:00	1:00	8:00	9:00

All sessions are live in Zoom or on Gather.town
All sessions in Zoom will be recorded and played on-demand for attendees

Day 1 29-Sep	Day 2 30-Sep	Day 3 1-Oct
Welcome "Hour"	**Symposia #1** (6:00-7:30) 4 x 15min 30min Q&A	**Poster spots** (5:30-6:00)
Select talks 1 4 x 15min 30min Q&A (7:15-8:45)	**BREAK**	**Poster Session #3** (6:00-7:30)
BREAK / 30 min student salons	**Networking 2** (7:45-8:15)	**BREAK**
Poster spots (9:15-9:45)	**BREAK**	**Select talks 5** 4 x 15min 30min Q&A (7:45-9:15)
Poster Session #1 (9:45-11:15)	**Select talks 3** 4 x 15min 30min Q&A (8:30-10:00)	**BREAK**
BREAK	**BREAK**	**BREAK / 45min student salons**
Networking 1 (11:30-12:30)	**Kavli Lecture** (10:15-11:15)	**BREAK**
BREAK	**Social hour**	**Symposia #2** (10:30-12:00) 4 x 15min 30min Q&A
Select talks 2 4 x 15min 30min Q&A (12:45-14:15)	**BREAK**	**BREAK**
BREAK	**Select talks 4** 4 x 15min 30min Q&A (12:30-14:00)	**Closing/ Awards** (12:15-12:45)
Panel Topics (14:30-15:30)	**BREAK**	
	Poster spots (14:15-14:45)	
	Poster Session #2 (14:45-16:15)	

Table of Contents

2021 Sponsors

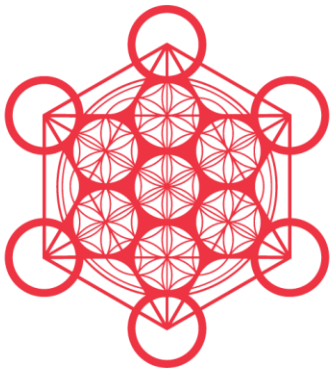

THE KAVLI FOUNDATION

Caltech
T&C Chen Center for Social and Decision Neuroscience

University of Zurich[UZH]

Wharton University of Pennsylvania
NEUROSCIENCE Initiative

NYU Institute for the Study of Decision Making

UNIVERSITÄT BONN
RHEINISCHE FRIEDRICH-WILHELMS-UNIVERSITÄT BONN

CENs
CENTER FOR ECONOMICS AND NEUROSCIENCE

Symposia Abstracts

Symposia #1 Variability in decision making: origins, mechanisms, and implications

S1.1 Imprecise learning drives variable but adaptive decisions under uncertainty in humans and artificial neural networks

Valentin Wyart[1]
[1]Ecole Normale Superieure

Making sense of uncertain environments constitutes a difficult yet ubiquitous challenge for human intelligence. Beside sensory errors and exploratory choices, recent research has identified imprecisions in the computations used to learn the structure of uncertain environments as a surprisingly large contributor to the variability of sensory- and reward-guided decisions. Taken together, the theoretical and experimental evidence scattered across psychology and neuroscience provides key insights into the origin, impact and function of this learning 'noise' for human decision-making. First, partitioning decision variability in terms of a bias-variance trade-off across sensory- and reward-guided decisions indicates that it is the limited computational precision of learning that drives the majority of decision errors under uncertainty. Second, neuroimaging data indicate that learning noise emerges from neural variability in the brain network activated during reinforcement learning and decision-making - a source of behavioral variability regulated by the noradrenergic neuromodulatory pathway. Finally, moving beyond the classical description of internal noise as a performance-limiting constraint on neural function and cognition, I will provide preliminary data that delineates the possible emergent benefits of computation noise for adaptive behavior in a range of adverse conditions.

S1.2 Correlated variability of stimulus and action codes tracks internal models for decision-making

Tobias Donner[1]
[1]University Medical Center Hamburg-Eppendorf

Objective: The mapping from states of the sensory environment (e.g., the orientation of a visual grating stimulus) to rewarded motor response (e.g., left- or right-hand button press) is governed by rules that can vary over time. A hallmark of cognition is the flexible adaptation of choice behavior to such rule changes [1]. This ability may result from the flexible updating of internal task models and the re-configuration of networks of neurons encoding stimulus features and motor actions [1, 2]. We set out to test the idea that this process might manifest in the correlation structure [3] of intrinsic fluctuations of neural population codes for sensation and action. Methods: Human participants performed a primary decision during fMRI, in which they discriminated the orientation (horizontal vs. vertical) of visual gratings in the visual field periphery and reported their choices with left- or right-hand motor actions. The rewarded stimulus-response mapping alternated after variable numbers of trials. The active rule was instructed explicitly in some blocks of trials ('instructed rule' task), but had to be inferred from a stream of ambiguous visual cues around fixation in other trial blocks ('inferred rule'). Participants' belief about the active rule for the latter (higher-order) decision was reconstructed by fitting a normative model [4] to participants' behavior in the primary decision. Intrinsic fluctuations in fMRI activity were quantified as the residuals after removing each voxel's response evoked by the primary decision. We

used multi-voxel pattern analysis to estimate time courses of orientation-selective activity in visual cortical regions and of action-selective activity in parietal and (pre-)motor cortical regions. Results: Participants' choices systematically followed the active rule in both, instructed and inferred rule tasks and were well-described by the normative algorithm for the inferred rule task. In both tasks, fluctuations of orientation-selective activity patterns and action-selective activity patterns were correlated. The sign of this selective correlation flipped along with the active rule. In the inferred rule task, the correlation tracked participants' graded belief about the rule, and the correlation was large when their choices followed the belief, but absent otherwise. Conclusions: Correlated variability of cortical activity tracks the dynamic re-configuration of networks for flexible decision-making and encodes internal models for decision-making. [1] Miller & Cohen (2001) [2] Shadlen & Kiani (2013) [3] Haefner et al. (2016) [4] Glaze et al. (2015)

S1.3 Exploring how humans explore: The contribution of catecholamines and mental health to distinct exploration strategies.

Tobias Hauser[1], Magda Dubois[1]
[1]University College London

Under which circumstances do we forgo good choice options to explore lesser known options? And how do we decide what to explore? Over recent years this exploration-exploitation trade-off has received much attention across different fields, from artificial intelligence to psychology. The appeal of near-optimal exploration algorithms from artificial intelligence has led to the proposal that humans use similarly sophisticated exploration strategies. However, such strategies are not scalable and quickly exceed human neurocognitive capacities. Here, I present recent and ongoing work demonstrating that humans deploy a variety of exploration strategies, including computationally much simpler exploration heuristic. I show how these different strategies induce variability in behaviour both within- and between-subjects. I present results from a double-blind placebo-controlled study investigating how specific catecholamines affect distinct exploration strategies. Moreover, I will present results from a pre-registered study investigating how differences in those exploration strategies are related to individual differences in personality and mental health traits. My talk will thus present a computational analysis of why humans are variable in their behaviour and how this within-subject variability is linked to inter-individual differences in behaviour and mental health.

S1.4 Metacognitive biases, but not model-based planning deficits, are improved following treatment with internet-based cognitive behavioural therapy.

Claire Gillan[1], Celine Fox[1], Chi Tak Lee[1], Tricia Xing Fang Seow[2], Kevin Lynch[1], Siobhan Harty[1], Richards Derek[3], Jorge Palacios[3], Veronica O'Keane[1], Klaas Enno Stephan[4], Claire Gillan[1]

[1]Trinity College Dublin, [2]UCL, [3]SilverCloud Health, [4]TNU Zurich

Prior research has identified alternations in both metacognitive processes and model-based planning in a range of mental health conditions. In order to translate these findings into mechanistic models and clinical impact, we need to move beyond cross-sectional designs. A key focus is the study of whether these cognitive deficits are stable, trait-like and difficult to alter, or fluctuate with symptoms over time and for example, following treatment success. The ongoing 'precision in psychiatry' (PIP) study was designed to answer these questions. PIP is a large transdiagnostic, 4-week observational study of low

intensity internet-based cognitive behavioural therapy (iCBT) targeting a range of mental health problems, but primarily anxiety and depression. To date, over 500 individuals have received iCBT through the SilverCloud platform and participated in our study. Self-reported clinical questionnaires, a perceptual decision-making task and a two-step decision making task were gathered at baseline and follow-up. Consistent with prior work, at baseline, individuals high in compulsivity exhibited excessive metacognitive bias (i.e. confidence) in their performance while those high in anxious-depression had reduced confidence. Also consistent with prior work, individuals high in compulsivity had model-based planning deficits. Following treatment, we observed significant improvements in both metacognitive bias (i.e. confidence in decisions) and model-based planning, but these can be partially explained by practice effects. Importantly though, clinical gains following treatment were associated with metacognitive improvements - that is, those who had the biggest reductions in anxious-depression following treatment, showed the biggest increases in confidence ($p<.01$). In contrast, there was no association between changes in anxious-depression or compulsivity and changes in model-based planning (both $p>.24$). These data suggest that metacognitive biases associated with anxious-depression are state-dependent and can be addressed through psychological intervention. In contrast, model-based planning deficits persist despite clinical gains, and have a more trait-like quality. These data add weight to the view that model-based deficits confer vulnerability for compulsive disorders and may be more difficult to modify than the metacognitive beliefs characteristic of anxious-depression.

Symposia #2 The Diverse and Critical Roles of Memory Retrieval Processes in Decision Making

S2.1 Episodic memory retrieval affects the build-up of decision variables in value-based choices

Peter Kraemer[1], Sebastian Gluth[2]

[1]University of Basel, [2]University of Hamburg

Objective In decision neuroscience, the concept of evidence accumulation (EA) has been established to model value-based choice formation. According to this framework, a decision maker integrates subjective value of different choice options over time until one option outweighs the other one by some margin. The mechanisms that define the input of EA signals, such as episodic memory retrieval, are usually seen as non-decisional processes that occur before EA starts. In the present study, we tested whether this assumption holds empirically. Methods In two sessions, participants (N=39) performed a remember-and-decide task, in which they made choices between money offers and snack items. Snack identity was presented either visually (value trials) or had to be recalled from memory (memory trials). In session 1, we obtained choices and response times (RT), and estimated a Drift Diffusion Model (DDM) in a hierarchical Bayesian framework. In session 2, we measured behavioral and EEG data to study the dynamics of EA on the cognitive and neural level. To identify relevant electrodes, we correlated simulated EA traces with time frequency representations. To this end, we compared the time courses of value and memory trials in these clusters using temporal permutation cluster tests. Results Behavioral data indicated comparable choice consistency between value and memory trials, accompanied with longer RT for the memory trials. Crucially, the estimated DDM suggested that RT differences cannot solely be attributed to non-decision time but also depend on lower drift-rate and larger boundary

separation. Hence, the DDM predicted that decision processes in memory trials depend on a longer build-up of EA signals. To test if the predictions of the DDM are supported by neural data, we identified electrode clusters which correlated with response-locked EA dynamics in the delta, theta, beta and gamma range ($p < .05$, Bonferroni-corrected). A time course analysis revealed higher frontal theta power for memory trials ($p < .001$) and a longer build-up time of frontal delta activity ($p = .046$). These results indicate that EA signals indeed emerged over longer time window in memory as compared to value trials. Conclusions Most EA models assume that memory retrieval can be attributed to non-decisional processes. While our overt behavioral data appeared to support this view, cognitive modeling and neural EEG responses suggest that decisions which involve episodic memory retrieval depend on a longer build-up of the decision variable. These results underline the relevance of considering the relative process dynamics of memory retrieval and valuation during choice formation.

S2.2 Process and content in decisions from memory

Wenjia Joyce Zhao[1], Russell Richie[2], Sudeep Bhatia[2]

[1]The Ohio State University, [2]University of Pennsylvania

Objective: Information stored in memory influences the formation of preferences and beliefs. The richness of this information, and the complexity inherent in interacting memory and decision processes, makes the model-driven analysis of memory-based decisions very difficult. In this paper we present a framework for building and testing formal models of naturalistic memory-based judgment and decision making. Our framework implements leading theories of memory search and decision making within a single integrated computational system, and can be used to characterize decision processes in everyday consumer, financial, health, ethical, legal, social, and policy decision tasks. Methods: Our modeling framework is inspired by a number of theoretical and technical insights. First, we use the thought listing experimental paradigm of query theory to jointly elicit memory data (in the form of recalled thoughts) and decision data. Second, we use the architectural assumptions of the context-retrieval and maintenance memory model to specify the interactions between memory and decision processes. With appropriate assumptions about the ways in which memory and decision making interact, we show that it is possible to decompose these processes into distinct components that can be separately fit to data. Finally, we use sentence vectors, obtained from recent deep-learning-based language models, to specify the semantic content of the thoughts that our participants retrieve from memory. We illustrate the power of our framework in three sets of experiments (2,484 participants) involving decision problems from a wide range of domains. Results and conclusions: In each experiment, we fit a total of 576 distinct memory-based decision models (each composed of a different combination of memory and decision models) to thought listing and choice data using Bayesian methods. Our fits reveal that memory-based decision making shows some domain-general regularities, including clustering effects in thought recall, context decay in both memory and decision making, and the accumulation of preferences to a threshold. Additionally, best-fit parameters (reflecting these regularities) predict key decision variables, such as the number of thoughts listed, the sequence of thoughts, and eventual decisions. They can also make out-of-sample predictions for the effects of experimental manipulations (e.g. thought primes) on decisions. These results showcase a powerful new framework that is capable of formally representing complex decision processes (and their individual-level and contextual correlates) in everyday decision tasks. Preprint: https://psyarxiv.com/nwxs2

S2.3 Experience replay supports planning and memory maintenance

Elliott Wimmer[1], Yunzhe Liu[1], Daniel McNamee[1], Raymond Dolan[1]

[1]University College London

Learning the structure of the environment and using our recent experiences supports our ability to plan and flexibly adapt to changes in the world. Flexible decision making may be supported by sampling of memories, potentially reflected in rapid sequential reactivation or 'replay' of past experiences. However, replay is also proposed to support the maintenance of memories themselves. Here, within the same experimental setting, we test these hypotheses in a model-based reward learning task where participants engaged with two separate, randomly alternating, environments. Using magnetoencephalography (MEG) and multivariate analysis, we find evidence for different replay signals during planning and following outcome receipt. During planning, replay strength for the current environment reflected demands for model-based planning and positively correlated with potential reward value. Following reward receipt, backward replay for the alternative environment increased as a function of decreasing recent experience, consistent with replay supporting a memory maintenance function. Our results provide novel support for two functions of memory replay, indicating that computational and on-going task demands modulate the degree to which awake replay of past experiences contributes to planning and memory maintenance.

S2.4 Mechanistic contributions of memory to decision-making impairments in Alzheimer's disease

Zhihao Zhang[1], Samira Maboudian[1], Ashley Jackson[2], Sang Ngo[2], Madhumitha Manivannan[1], Andrew Kayser[2], Ming Hsu[1], Winston Chiong[1]

[1]University of California, Berkeley, [2]University of California, San Francisco

Life is not a multiple-choice test. Because real-world decisions are often open-ended, with choice options conceived by decision-makers themselves, the quality of decisions may heavily rely on the retrieval of options, as failure to generate promising options limits, and often eliminates, the opportunity for choosing them (Zhang et al., 2021, PNAS). Because of the demand on memory retrieval, such decisions may pose substantial challenges to patients with Alzheimer's disease (AD), the hallmark of which is substantial and worsening memory deficits. Our sample included AD patients (N = 13), along with healthy controls (HC; N = 16) and behavioral-variant frontotemporal dementia (bvFTD) patients (N = 13) as a dementia control group without memory impairments. The experiment consisted of (1) a semantic fluency (SF) condition, where subjects recalled as many items as possible from a category, (2) an internal-menu choice (IMC) condition, where subjects chose their two most preferred items from a category without a menu, and (3) an external-menu choice (EMC) condition, where subjects made the same choice with a predefined menu. 14 categories covering a variety of goods were used, and subjects performed each condition for all categories. The order of SF and IMC was counterbalanced and EMC was always the last. Of primary interest is choice switches between IMC and EMC in the same category, with the rationale being that memory constraints in IMC would be lifted in EMC due to the provision of a menu. By examining the occurrence and nature of such switches and linking them with performances in SF, we could characterize the extent and causes of decision impairments in different groups. As expected, AD patients recalled fewer items in SF than both HC ($p < 0.005$) and bvFTD groups ($p < 0.01$) across categories. Importantly, this difference in memory retrieval capabilities were related to IMC to EMC choice switches: the AD group made more switches than both bvFTD ($p < 0.05$) and HC ($p < 0.01$), and a significant correlation was observed across all groups between the average number of items

recalled in SF and the average choice switches. In parallel, at the item level, we also found that subjects tended to switch to more difficult to recall items. Furthermore, AD patients made higher memory-driven switches than HC, while bvFTD patients showed more value-driven switches. These findings provide novel evidence for an overlooked source of decision impairments in AD and connect models of open-ended decisions with longstanding findings of memory impairments in AD. They bear implications for new interventions that help improve decision quality in AD and related dementia.

Oral Sessions Abstracts

Oral Session #1

O1.1 Attention constraints and learning in categories

Rahul Bhui[1], Peiran Jiao[2]

[1]MIT, [2]Maastricht University

When different stimuli belong to the same category, learning about their attributes should be guided by this categorical structure. In particular, attention constraints should rationally lead people to focus on learning about shared qualities and neglect individual differences. This mechanism has been proposed to explain why financial investors often focus on industry-level information while neglecting idiosyncratic stock-specific data. Although prominent perspectives in psychology, neuroscience, and economics maintain that attention is spent where it has maximum benefit, it remains unclear to what extent attention allocation is guided (and hence biased) by categorical structure. Here, we experimentally tested whether biases toward learning at the category level may be driven by the optimal allocation of attention. We conducted three preregistered experiments (total N = 438) using an information sampling paradigm with mousetracking, in which participants had to estimate the values of hypothetical stocks based on streams of incoming information about individual stocks and their common industry. Participants could choose which information to acquire moment-to-moment, and we measured the amount of time allocated to acquiring information about each variable. If attention is deployed optimally, one's focus on shared qualities should adjust flexibly to environmental features that affect the value of information attainable from each source. This core idea yields several predictions. First, when members of a category are similar, there is little to be gained by learning each one's unique qualities; thus, people should focus more on information at the category level when idiosyncratic variation is low relative to shared variation. Second, continuing to accumulate information about a given variable yields diminishing returns; thus, people should focus more on information at the category level when they face severe attention constraints. Finally, category-level information reduces uncertainty about every member, and so its value scales with category size; thus, people should focus more on information at the category level when the category contains many members. We found support for each of these predictions: participants spent more time acquiring industry-level information when idiosyncratic stock variation was low, when time constraints were more severe, and when the industry contained more stocks. Our results thus provide evidence that categorical biases can be driven by a rational response to attention constraints.

O1.2 Foraging vs Value-Comparison Reinforcement Learning Models of Sequential Decision-Making

Dameon Harrell[1], Cathy Chen[1], Collin Meyer[1], David Darrow[1], Nicola Grissom[1], Becket Ebitz[2], Alexander Herman[1]

[1]University of Minnesota, [2]University of Montreal

OBJECTIVE: Some decisions must be made with limited information about the available options. This limitation can lead to an explore (choose a new option) vs exploit (stay with the current one) dilemma. Reinforcement learning (RL) is traditionally used to model sequential decision-making. Typically, RL models implement action selection via value comparison. In environments with few options, comparisons are simple. However, as the number of options increases, so does the complexity of a value-comparison strategy. Moreover, the animal literature suggests an alternate decision-making strategy in the foraging (FG) model. Instead of value comparison, the FG model uses a compare-to-threshold strategy which is less sensitive to the number of options in an environment. Whether human sequential decision-making is better described by an RL or an FG model is unknown. METHODS: Discerning which of these strategies might best explain observed behavior can be investigated by quantifying how the ratio of exploration to exploitation changes as a function of environmental richness. Here, environmental richness is defined as the total value of all available options, where value is the probability of reward. RL models, due to their value-comparison policies, predict increased exploration in both extremely poor and extremely rich environments compared to intermediate ones, as there can only be small differences between option values at the extremes of environmental richness. Alternately, an FG model predicts the proportion of exploration should decrease monotonically as richness increases. To test these predictions, we compared the two models using a restless multi-armed bandit task. In addition, we collected data (online) from 220 human participants performing the same bandit task. Classification of exploration and exploitation was determined by modeling decisions as the consequence of latent states in a hidden Markov model. RESULTS: The FG and the RL model performed differently in rich environments. In plausible parameter ranges, the RL model produced increased rates of exploration at the extremes of richness, whereas the FG model exhibited a monotonic decrease in exploration as levels of richness increased. Furthermore, we found that the average behavior of the human participants was better captured by the FG model. Specifically, the proportion of exploration decreased monotonically with increased environmental richness, qualitatively similar to the FG model but not the RL model. CONCLUSIONS: Our results suggest that human decision making is sensitive to environmental richness and that this sensitivity is best explained by a compare-to-threshold mechanism.

O1.3 Prefrontal-hippocampal involvements in distinct decision strategies during human foraging behavior for real liquid rewards.

Reiko Shintaki[1], Daiki Tanaka[1], Shinsuke Suzuki[2], Takaaki Yoshimoto[3], Norihiro Sadato[3], Junichi Chikazoe[3], Koji Jimura[1]

[1]Keio University, [2]University of Melbourne, [3]National Institute of Physiological Sciences

Objective: Foraging is a fundamental food-seeking behavior in a wide range of species involving continuous decision-making regarding whether to remain in the current environment or move on to a new environment. This decision contrasts with intertemporal choice, in which a one-shot decision is made, limiting empirical evidence about whether the behavioral agent would continue waiting after the

decision. In this study, we sought to overcome this limitation using a foraging task for real delayed liquid rewards, and to examine the value of anticipation and underlying neural mechanisms (refs. 1-3). Methods: Human participants performed a foraging task while being subjected to high spatiotemporal fMRI. In each set of trials, participants were presented with a novel rewarding environment, and directly received a real liquid reward delayed by tens of seconds (experience trial; ref. 3). They were then presented with the same environment, and again waited for a reward (foraging trial). Importantly, in contrast to intertemporal choice tasks, participants were unsure about when they would receive a reward. In exchange, they were able to stop waiting for the delayed reward and to move on to the next experience trial with a novel environment, whenever they preferred. Results and Discussion: Survival analysis of the foraging trial revealed that participants stopped waiting more frequently and sooner after they experienced longer delays and received smaller rewards in the experience trial. While participants anticipated a future reward in the foraging trial, temporal dynamics of anticipatory utility (AU; ref. 1) was modeled by Bayesian learning of probability density of expectation for the reward (refs. 2, 3). Whole brain exploratory neuroimaging analysis revealed that the AU dynamics was associated with the anterior prefrontal cortex (aPFC) and hippocampus (HPC). Interestingly, when the aPFC signal was attenuated during the delay period, participants stopped waiting. In aPFC, the AU dynamics were more prominent in individuals with exploratory strategies [smaller area under the curve (AuC) in the survival function], whereas individuals with staying strategies (greater AuC) exhibited more prominent dynamics in HPC. In the experience trial, the AU dynamics were associated with aPFC and HPC and reflected choice strategies similar to those in the foraging trial. These results suggest that exploratory and staying strategies involve aPFC and HPC, respectively, highlighting prefrontal and hippocampal engagements in human foraging behavior. References: (1) Loewenstein 1987. Econ J 97:666. (2) Jimura et al. 2013. J Neurosci 33:344. (3) Tanaka et al. 2020. J Neurosci 40:9736

O1.4 Trait somatic anxiety is associated with reduced exploration and underestimation of relative uncertainty

Haoxue Fan[1], Samuel Gershman[1], Elizabeth Phelps[1]

[1]Harvard University

People constantly face the "explore-exploit" dilemma: should I stick with the current best option (exploit), or should I try something else that could potentially be better (explore)? Though reduced physical exploratory behavior has been long assumed in animal anxiety models, both positive and negative relationships between anxiety and exploration during decision-making in humans have been documented. One reason for the past mixed results is the multidimensionality of exploration and anxiety. In the current study, we examined how people explore under uncertainty and how trait anxiety components interact with the exploration process. To quantify different exploration strategies, we used a two-arm bandit task and manipulated different kinds of uncertainty by changing the volatility of each arm. Across two well-powered studies (study1: N = 501; study2: N = 483), We demonstrated that people use a hybrid of directed, random, and softmax exploration strategies, which are respectively sensitive to relative uncertainty, total uncertainty, and relative value. In terms of trait anxiety, people high on somatic trait anxiety rely less on directed exploration i.e., they are less likely to choose the option with high relative uncertainty (study1: ß = -0.07, t(150285) = -2.83, p = .004, CI = [-0.11, -0.02]; study2: ß = -0.05 , t(144885) = -2.95, p = .003, CI = [-0.09, -0.02]). Somatic trait anxiety is additionally associated with reduced softmax exploration, manifesting as lower choice stochasticity(study1: ß = 0.12 , t(150285) =

2.63, p = .008, CI = [0.03, 0.21]; study2: ß< = 0.14, t(144885) = 3.59, p < .001, CI = [0.06, 0.21]). In study2, we further tested whether somatic trait anxiety is related to biased uncertainty and/or value computation. At the end of each block of the two-armed bandit task, participants were asked to report their reward predictions of two bandits and their confidence in the estimates. By comparing participants' subjective estimates to normative estimates derived from a Kalman Filter, we found that people high on somatic trait anxiety underestimate the relative uncertainty in the environment (ß = 0.2 , t(12052) = 4.00, p < .001 , CI = [0.10, 0.30]) while somatic trait anxiety doesn't influence relative value estimation (ß = -0.54, t(12052) = -1.43, p = .15, CI = [-1.27 0.20]). Together, these results bring insights into the relationship between anxiety and exploration in human decision-making, showing a selective impact of trait somatic anxiety on directed and softmax exploration as well as perceived relative uncertainty.

Oral Session #2

O2.1 Efficient coding accounts for faster and more accurate choices on high-valued items

Seungji Lee[1], Sung-Phil Kim[1],

[1]Ulsan National Institute of Science and Technology

One of the goals of decision-making research is to identify cognitive processes underlying choice behavior similar to the ones we make in our lives as consumers. A common observation in the choice of goods is that consumers tend to select a set of more valuable goods and choose the one within the set. However, it has not been fully understood how value-based choice behavior varies with the value level (i.e., high or low) of options. In this study, we aim to investigate subject-value based choice behavior with various value levels in terms of the speed and accuracy of choice at different levels of choice difficulty (i.e., smaller or larger value difference). We also aim to account for an underlying value computation process by comparing two alternatives of mechanisms - efficient coding and attention - in the framework of sequential sampling models (SSMs) for binary choice. Here, we chose snacks as a target goods. To categorize snacks into different value levels, we first collected a set of snacks that participants had tried before and evaluated how much they want to eat each collected snack by measuring subjective ratings from -10 (not at all) to 10 (very much) for three times and the mean ratings were used for classifying snacks into three value levels (i.e., low, medium, and high). Each snack could belong to only one value level or none. Choice pairs were created using two snacks in the same value level with four levels of value difference. Accordingly, there were 12 conditions of choice pairs with a maximum of 30 trials for each condition. The data of 24 participants (12 female, age: 24.13±3.41 years old) were analyzed. Behavior results showed that the larger the value difference was, the faster the choice was made with higher accuracy. Increases in speed and accuracy as a function of the value difference were steeper when choosing between high-valued than low-valued items. Based on the behavioral results, we simulated data using three SSMs and compared them to the behavioral data: a typical SSM only concerning choice difficulty (tSSM); an SSM employing efficient coding for high-valued items (eSSM); and an SSM employing attention to high-valued items (aSSM). From the model analysis, eSSM was found to be the only model that could describe the interaction between the value level and the value difference on choice behavior. Our findings highlight the role of efficient and precise valuation in the choice of daily goods and help to understand subjective value-based decision-making processing.

O2.2 Irrational choice via curvilinear value geometry in ventromedial prefrontal cortex

Katarzyna Jurewicz[1], Brianna Sleezer[2], Priyanka Mehta[2], Benjamin Hayden[2], Becket Ebitz[1]

[1]Université de Montréal, [2]University of Minnesota

OBJECTIVE: Many patterns of irrational choice have found satisfying explanation in algorithmic or computational bounds on decision-making. Though others still defy psychological explanation, they could be due to bounds on hardware: to constraints on the brains' ability to faithfully represent the value of different offers. Yet, because we have only recently developed methods to study the geometry of neural representations, we still do not know whether there are constraints on the neural representations of value nor what impact these may have on choice. METHODS: Two rhesus monkeys performed a value-based decision-making task while we recorded in the ventromedial prefrontal cortex (vmPFC; area 14; 122 isolated neurons), a prefrontal region implicated in economic decision-making. The monkeys repeatedly chose from random sets of sequentially-presented options whose appearance clearly cued some volume of juice. The monkeys were highly accurate: they chose the best option in each set 94% and 96% of the time, respectively. Neural state-space analyses were then used to probe the geometry of the population-level offer value representations during offer viewing periods. RESULTS: Offer value representations were precisely and predictably arranged in vmPFC, but also nonlinearly. Instead of following a straight line, offer values traced a curved path through neuron-dimensional space (significantly more curvature than linearized control, $p < 0.0001$, bootstrapped test). Because decoding from a curvilinear manifold is a challenge whose difficulty scales with the range of encoded items, the curvature predicted a surprising violation of rational choice theory: the idea that terrible offers--offers that will never be chosen--should not affect preference for better offers. Indeed, monkeys were much more likely to confuse good options when an irrelevant offer was worse (and the range was bigger), compared to when an irrelevant offer was better (17% change in accuracy; significant in both monkeys, $p < 10\text{-}12$). This meant that both monkeys made the best decisions when the offers were closest together, but the warping due to curvature was minimized. CONCLUSIONS: Theoretical work on value-based decision-making often assumes that values are linearly represented within the brain, a geometry that is likely critical for generating rational economic decisions. However, we find that (1) value is represented nonlinearly in at least one economic decision-making region, and (2) behavior violates rational choice theory in exactly the way we expect from this curvilinear geometry. Neural curvature could contribute to systematic patterns of irrational choice.

O2.3 Neural codes in early sensory areas maximize fitness

Jonathan Schaffner[1], Philippe Tobler[1], Todd Hare[1], Rafael Polania[2]

[1]University of Zurich, [2]ETH

Objective: It has generally been presumed that sensory information encoded by a nervous system should be as accurate as its biological limitations allow. However, accurate representations of sensory signals do not necessarily maximize an organism's chances of survival. Thus, organisms should use fitness-maximizing encoding schemes in many cases. To test this hypothesis, we developed a unified normative framework for fitness-maximizing encoding by combining theoretical insights from neuroscience, computer science, and economics. Methods: We tested the predictions of the model using responses from blowfly retinal neurons and novel experiments in humans (n=25 in Experiment 1,

n=30 in Experiment 2). In experiment 1, participants trained on a perceptual task in one of two conditions that used the same prior distribution of stimuli. In the accuracy condition (Kacc), they were rewarded with a fixed amount whenever they correctly indicated which of the two Gabor patches was oriented closer to 45-degrees. In the reward condition (Krew), participants were trained on the same task, but the reward was determined by the orientation they selected (e.g., choosing a 40-degree patch over a 38-degree patch yielded 40 points). Critically, in an information-maximizing coding scheme (infomax) will maximize reward in Kacc, but a fitness-maximizing coding scheme (fitmax) will maximize reward in Krew. We measured participants orientation estimation accuracy and repulsive bias before and after they trained in each context. For experiment 2, we trained different participants in Krew only. However, we tested the participants' estimation accuracy in the retinotopically-specific trained location and other untrained locations. Thus, experiment 2 allowed us to test if the training influenced spatially selective neurons early in the visual pathway. Results: In the blowfly, we found that neural codes that maximize reward expectation--and not accurate sensory representations--best account for retinal LMC activity. The results of the experiments in humans were also consistent with the predictions made by a fitmax code. For example, decreases in repulsive biases were greater following Krew than Kacc (interaction K∗time(after- before): P = 0.0085). Also, we found that reductions in estimation bias were present in trained (P = 0.0078), but not untrained locations (P = 0.43; interaction location*time: P = 0.0425). Conclusions: We showed that the earliest stages of sensory processing have evolved to encode environmental stimuli to promote fitness maximization, and not necessarily to maximize perceptual accuracy.

Oral Session #3

O3.1 Sources of confidence in value-based choice

Jeroen Brus[1], Marcus Grueschow[1], Rafael Polania[1]

[1]ETH

Introduction "Which meal would you like, chicken or pasta? Chicken please. . hm not sure. No sorry, I prefer pasta". To evaluate our decisions without feedback is a fundamental aspect of cognition, which can be used to revise decisions and guide future behavior. This subjective evaluation is known as confidence, which is highly relevant for economic decisions. Objective Confidence is statistically defined as the chance that a decision is correct. A central question is whether empirical confidence reflects statistical confidence and accurately represents decision uncertainty. Here, the goal is to study the sources of confidence in value-based decisions. Methods The experiment consists of a food rating task that was repeated, which provides a measure of variability in the value estimates, and an incentive-compatible task between the rated food items. Subjects (n=33) also indicate their confidence in their choices. During the incentive-compatible task the eyes of the participant were tracked. We performed a model free analysis and implemented a recent evidence accumulation model that takes into account attentional effort: the Gaze-weighted Linear Accumulator Model. Typically attentional effort (θ) is assumed to be agent-specific. Instead, we allow θ to fluctuate from trial-to-trial. We compare two generative confidence models: a previously proposed heuristic model, based on the location of the loser accumulator, and a normative confidence model based on the statistical definition. To study co-fluctuations of confidence and decision model parameters, we adopt a joint modeling approach. Lastly

we fitted an efficient coding model that uses the statistics of the environment to dissociate between encoding and comparison noise. Results Surprisingly, our model free analysis showed that fluctuations in the precision of value encoding have no influence on confidence reports. Model comparison shows that confidence is best explained by the normative model that allows θ to fluctuate (Bayes Factor > 1000). Higher levels of θ lead to higher confidence reports ($P < 0.001$). We replicated all results above using an independent data set. Our efficient coding model showed that confidence emerges from comparison noise instead of encoding noise. Conclusions Our model free analysis suggests that confidence is not influenced by uncertainty in the input values. Instead, the generative models of confidence show that attentional effort fluctuates and participants incorporate this in their confidence reports. Our efficient coding model results converge with the model free analysis and show that encoding noise does not influence confidence reports, but comparison noise does.

O3.2 Goal-dependent memories in value-based choice

Mariana Zurita[1], Benedetto De Martino[1]

[1]University College London

We make value-based decisions every day. These decisions are usually based on values or preferences shaped by our past experiences (e.g. a favourite walk home from work). However, if the main goal or the context changes (e.g. tiredness), the way we evaluate our choices may also change. In a recent study, we found that value representations in the brain are re-shaped by different goals (Castagnetti, Zurita & De Martino 2020). In this study, we investigated how goals shape the way memories are constructed and used under different behavioural demands. We devised a behavioural task in which subjects navigated a 2D maze to find a specific token (i.e. task goal) at unknown locations. Different items were scattered in the maze across the possible paths and provided cues to the location of the token. We investigated how subjects changed their representation of the environment through direct memory tests and by modelling their spatial representation. This setup allowed us to test whether memory traces were influenced by the behavioural goal, affecting the recall of (1) items that are closer to the current goal, and (2) items that are more informative (i.e. have a higher prediction error) to the current goal. To have a better understanding of the computational process underpinning this behaviour, we also tested how participants' spatial representation had been affected by different goal-values, using a reinforcement learning algorithm called successor representation (SR). The SR is a temporal-difference algorithm that incorporates complex representations of the environment within its value calculations. It has been shown to be a good model for the predictive function of the hippocampus (including grid/place cells). We hypothesised that in spite of being under the same objective space, subjects would build a deformed memory representation under each goal (changing the representation in the SR algorithm), a process guided by their goal-value policy. This work aims to demonstrate how goals shape value and, in turn, foster flexible behaviour by modifying the internal representations of the environment stored in memory.

O3.3 Contrasting range normalization and divisive normalization in human reinforcement learning

Sophie Bavard[1], Stefano Palminteri[1]

[1]INSERM U960

Context-dependent learning has been shown to lead to irrational choices in humans. This is specifically true when the options are extrapolated from their original learning context (Palminteri et al. 2015; Klein et al. 2017; Bavard et al. 2018, 2021). In a previous study, we showed that this process was well captured by a dynamical range normalization model, inspired by the range-frequency theory (Parducci 1995) and electrophysiological findings in monkey (Padoa-Schioppa 2009). However, the two-armed bandit previously used is ill-suited to precisely characterize the functional form of context-dependence as range normalization or divisive normalization (Louie et al. 2013; Webb et al. 2020). To fill this gap, we designed a new online-based learning task simultaneously manipulating the number of options per context (2-armed bandit versus 3-armed bandit) and the range magnitude of the options, by varying their expected values. We also included an explicit valuation phase where participants had to report their estimation of each option. Behavioral and computational analyses seriously challenge divisive normalization, but suggest that simple range normalization cannot account for all behavioral patterns. Together, these results shed new light on the mechanisms of context-dependent learning in humans.

O3.4 The trembling hand unraveled - the motor dynamics and neuronal sources of choice inconsistency

Vered Kurtz-David[1], Adam Hakin[2], Asaf Madar[2], Noa Palmon[2], Dino Levy[2]

[1]New York University, [2]Tel Aviv University

Objective: The "trembling hand error" is a synonym for accidental outcomes in decision-making. Nonetheless, as the term implies, such mistakes could result from a violation of utility maximization, from motor tremor that accompanied choice execution, or both. We argue that choice inconsistency stems from noisy neural dynamics, and aim to differentiate between two possible sources of noise: noisy value computations originating in value related brain areas, as opposed to noise during motor execution originating in motor-related brain regions. To do so, we examine the dynamics leading up to choices by analyzing mouse trajectories and neural activity. Methods: We conducted a behavioral study (n= 89) and a neuroimaging study (n=42). Subjects completed a value-based risky-choice task using budget sets to measure their inconsistency level on a given trial. On each trial, in addition to subjects' choice, we track their mouse movements. We extract 37 features from the recorded trajectories to characterize choice dynamics that underlie either motor or value sources of inconsistency levels. Subjects also completed two supplementary motor tasks that were designed to examine pathways without value modulation. We investigate which mouse features correlate with inconsistency levels to identify typologies of motor execution that are related to choice inconsistency. We also compare mouse trajectories from the risky-choice task to trajectories in the two supplementary tasks, to differentiate between motor and value-related dynamics. In the neuroimaging study, we also identify regions that correlate with choice inconsistency, and systematically look for neural representations of mouse features across motor and value brain areas. Importantly, we compare the representations of mouse features with and without value modulation. Last, we relate spatial and temporal patterns in these regions to inconsistency levels and dynamics identified by the mouse features. Results: We find that choice dynamics not related to value computation account for some of the variation in inconsistency levels. We further replicate our previous findings, suggesting that choice inconsistency correlates with value modulation in the brain. Moreover, we find that choice dynamics have neural footprints in M1 and SMA, but are also strongly represented in value-related brain areas. Conclusions: These findings imply

that certain elements in motor execution are related to choice-inconsistency, and further indicate that mouse tracking can reveal the underpinnings of choice. The neuroimaging results suggest that patterns of both value and motor cues contribute to inconsistency levels.

Oral Session #4

O4.1 Computational modelling of learning and action initiation from childhood to adolescence

Ruth Pauli[1], Inti Brazil[2], Gregor Kohls[3], Stephane De Brito[1], The FemNAT-CD consortium /[1], Patricia Lockwood[1]

[1]University of Birmingham, [2]Donders Institute for Brain, Cognition and Behaviour, Radboud University, [3]University Hospital RWTH Aachen & TU Dresden

Objectives: The ability to learn from reward and punishment is an essential aspect of development across childhood and adolescence (Nussenbaum & Hartley, 2019). New advances in computational modelling can precisely quantify learning and have suggested developmental differences, with adolescents learning more strongly from reward than punishment compared to adults (Palminteri et al., 2016), and showing an attenuated 'go' bias to initiate actions (Raab & Hartley, 2020). However, most work to date has used small sample sizes and/or adolescent samples based in single countries. Here we used a computational modelling approach to investigate changes in reward and punishment learning and action initiation from childhood to adolescence, using a large sample (N = 743, aged 9-18 years) recruited from 13 European countries (Freitag et al., 2018). Participants completed a learning task in which they were required to learn by trial and error whether to press a button in response to an abstract shape to obtain rewards (points), or withhold pressing to avoid punishment. Each shape was associated with a different gain or loss magnitude, and outcomes were deterministic (Kohls et al., 2020). We fitted and compared a set of computational models using a hierarchical expectation maximisation fitting procedure and Bayesian model comparison approach. Results: Model comparison demonstrated that learning was best captured by a model with a single temperature parameter, separate learning rates for reward and punishment, a constant 'go' bias, and an outcome magnitude sensitivity parameter. Punishment learning rate increased with age (robust linear mixed effects regression: β=-0.0002, t=3.81, p<0.001), whereas the bias to 'go' decreased with age (β=-0.001, t=-4.68, p<0.001). By contrast, reward learning rates appeared fairly stable between ages 9 and 18 years (β=0.00005, t=1.08, p=0.39). Bayesian correlations confirmed strong evidence for a correlation between age and punishment learning rate (BF10=1006.24), and between age and 'go' bias (BF10=5062.17), but weak or no evidence for a correlation between age and reward learning rate (BF10=0.08). Conclusions: These results suggest that there are important developmental differences in learning and action initiation from childhood to adolescence. Crucially, a bias to perform an action declines across adolescence, whereas learning rates from reward remain stable. Previous demonstrations of age-related changes in reward learning could therefore in part reflect changes in 'go' bias rather than reward learning per se. In addition, punishment sensitivity and action initiation are subject to developmental changes in normative populations.

O4.2 The Neural Mechanisms of Affect Misattribution on Judgment and Evaluation

Aiqing Ling[1], Nathalie George[2], Baba Shiv[3], Tobias Kalenscher[4], Hilke Plassmann[5]

[1]UCD Michael Smurfit Graduate Business School, [2]Brain Spine Institute (ICM), Sorbonne University, [3]Stanford Graduate School of Business, Stanford University, [4]Institute of Experimental Psychology, Heinrich Heine University Düsseldorf, [5]INSEAD

Why do we think our lives are better on sunny vs. rainy days? Previous work has focussed on how good sunny weather makes us feel, i.e. its valence when colouring our judgements while the potential interaction of valence and arousal of affective states has been under-researched during such affect infusion processes. We demonstrated in a series of studies combining fMRI, skin conductance & facial affect recording and behavioural approaches for the first evidence for the novel arousal transport hypothesis (ATH): The infusion of incidental affect necessitates arousal to transfer affective valence on subsequent judgements. We then showed in two natural experiments that when arousal is selectively dampened due to a stressor - terror attacks in S3 and COVID19 in S4 - affect infusion processes are attenuated. S1 (N=20) & S2 (N=47):participants evaluated their enjoyment of wines / images after randomly winning money or not from an incentive compatible lottery. We found that how much participants enjoyed the wines / images was mediated by the neural activity of valence encoding regions (S1: β=.0002, 95% CI [.0001 .0002]) / valence encoded in facial affect (S2: β=.02, 95% CI [.01 .04]). Further, the neural activity of arousal encoding regions (S1: βhigh vs. low=.002, 95% CI [.0006 .003]) / skin conductance responses (SCRs) (S2: βhigh vs. low=.04, 95% CI [.003 .09]) positively moderated the strength of the mediating effects. S3 (N=92):a group of participants performed the same evaluation task as in S2 up to 7 days after terror events (stress group, SG) and a group of participants with similar demographics did the study in the absence of a terror attack (control group, CG). We found intact facial valence encoding in response to affective images / monetary gains in both CG (βimg=.026, p<.01; βreward=.025, p<.01) and SG (βimg=.016, p=.03; βreward=.022, p<.01). However, SCRs towards these stimuli were only found in CG (βimg=.05, p<.01; βreward=.028, p<.01) but not in SG (βimg=.0015, p=.55; βreward=.01, p=.11). This evidence indicated a selective bluntness of arousal due to the stressor. As a result, the mediating effect of facial valence was insignificant for SG participants (β=.01, 95% CI [-.007 .04]). S4 (N=547):a group of Chinese and French participants performed the same image evaluation task as in S2 during the COVID19 pandemic. We measured participants' self-reported stress and sampled their hair to analyze stress hormones. We found that chronic stress scores negatively moderated the extent of affect infusion (β=-.023, p<.01), echoing the idea that the effects of stress aftermath (vs. acute stress) may down-regulate arousal responsiveness.

O4.3 Sensitivity to contextual effects during reinforcement learning in human addiction

Maëlle Gueguen[1], Hernan Anlló[2], Darla Bonagura[1], Julia Kong[1], Sahar Hafezi[1], Stefano Palminteri[3], Anna Konova[1]

[1]Rutgers University, [2]Waseda University, [3]Inserm - ENS Paris

Recent work suggests human decision-making behavior is highly context-dependent. Rather than evaluated on an absolute scale, value is influenced by the context in which it is computed and the range of available reward. This eases the identification of the highest/lowest valued options but is only valid in a given context and can lead to choice errors when the context changes. In drug addiction, the reinforcing properties of the drug and withdrawal can have outsized effects on the effective "range" of experienced reward, potentially altering sensitivity to context. Here, we explore how people with chronic drug addiction adapt to contextual changes in reinforcement learning, and whether they rely more on an absolute value-based strategy or a context-dependent one, and how. 28 individuals with

opioid use disorder (OUD) (mean [SE] age, 48.8 [2.3] years; 10F) and 28 matched controls (52 [2.5] years; 13F) completed a three-phase task designed to produce robust context effects. In the first phase, participants learned to choose the best option in 4 pairs of cues associated with different magnitude and probability of reward [expected values (EV)]. In the transfer phase, the same cues were rearranged to create new pairs, such that the correct options during learning were not necessarily correct during transfer. Finally, we probed choice based on the explicitly stated EV difference between cues (ΔEV) matching the transfer pairs. This design allowed us to quantify choice during acquisition (learning), extrapolation (transfer) and estimation (explicit) of EV. All participants identified the higher EV cues in the learning phase, with no group difference. In line with context-dependent choice, performance during transfer dropped below chance level for the pair for which previous cue rank conflicted with current ΔEV (one-sample t-test vs. chance across the sample, p=0.04) and was significantly below that observed in the explicit phase (p=0.03). By contrast, a boost in performance was observed in OUD relative to controls for the pair for which previous cue rank and current ΔEV aligned (two-sample t-test, p=0.004). This boost allowed performance in OUD to reach levels observed in the explicit phase when ΔEV could be easily estimated (paired t-test, p=0.20) but was dampened by longer duration (Spearman ρ covarying for age, p=0.03) and more recent opioid use (p=0.002). Overall, people with drug addiction may be more sensitive to contextual modulation of value than their healthy counterparts. Drug use seems to diminish the benefit of a context-dependent strategy when previous rank and current ΔEV align but does not affect the decision process when the two are in conflict.

O4.4 Effects of age and medial temporal lobe atrophy on memory-guided decision-making

Karolina Lempert[1], Michael Cohen[1], Kameron MacNear[2], Frances Reckers[1], Laura Zaneski[1], David Wolk[3], Joseph Kable[1]

[1]University of Pennsylvania, [2]University of Illinois ? Urbana-Champaign, [3]Hospital of the University of Pennsylvania

OBJECTIVE: From choosing whether to buy a product again to choosing whether to respond to a familiar but suspicious e-mail, people often draw on memories of single previous episodes to make decisions. Research in young adults has shown that retrieving an accurate association between a stimulus and its value (i.e., associative memory) is necessary to appropriately approach high-value stimuli and avoid low-value stimuli. Given that older adults show a decline in associative memory, here we examined the effects of aging on memory-guided decision-making, in social and non-social domains. We also examined links between atrophy in the medial temporal lobe (MTL), a brain structure critical for episodic memory, and memory-guided decision-making. METHODS: Young (n=50; aged 18-35), middle-aged (n=73; aged 36-59) and cognitively normal older (n=89; aged 60-92) adults studied a series of faces along with how much the people pictured shared in a dictator game previously ($5 or $0 of $10). Participants also saw house images representing lotteries worth $5 or $0. Then participants made real-stakes choices about whether to interact with each person or house again. This was followed by recognition and associative memory tests. A subset of older adults (n=51) underwent structural neuroimaging to obtain volume/thickness measures of MTL subregions, including hippocampus, entorhinal cortex, parahippocampal cortex, and perirhinal cortex regions BA35 and BA36. RESULTS: Replicating previous research, correct associative memory for a stimulus and its value was necessary for making optimal decisions, in all age groups. Age was associated with worse associative memory ($r=-0.23$; $p<0.001$), and worse decision-making ($r=-0.46$; $p<0.001$). In fact, older adults were not above chance at

approaching high-value stimuli and avoiding low-value stimuli. Especially in the social domain, this decision-making deficit was not driven solely by poor associative memory: older adults simply did not avoid people that they confidently remembered as being unfair. The only MTL region in which median cortical thickness was associated with memory-guided decision-making was BA35 (r=0.33; p=0.02), the first deposition site of Alzheimer's neurofibrillary tangles. CONCLUSIONS: Poor associative memory combined with a bias toward approaching familiar faces (regardless of remembered value) might make older adults vulnerable to exploitation. The finding that BA35 atrophy (which may distinguish healthy from pathological aging) was linked with decision-making even in cognitively normal older adults suggests that subtle changes in the ability to use memory to guide choice may signal underlying Alzheimer's pathology.

Oral Session #5

O5.1 Computational mechanisms used for coordination learning in free ranging baboons (Papio papio)

Toan Nong[1], Remi Philippe[1], Nicolas Claidière[1], Jean-Claude Dreher[1]

[1]CNRS

One important question is to understand the algorithms developed by the brain to cooperate with conspecifics. Here we studied free ranging interacting baboons. We hypothesized that social rank is a key factor to explain coordination between individuals. We tested and compared a number of computational models accounting for coordination learning, such as Bayesian Learning, Reinforcement Learning, heuristics, mentalizing models and others. We used a simple two-by-two coordination game to test baboons (N=7) either interacting with a conspecific in a social condition or non-interacting in a control non-social condition. They have the choice between two different targets and are rewarded if they coordinate on the same choice. In fact, the actual correct choice was given by a cooperative algorithm, the same one in both conditions. The specificity of our experimental setup is that our baboons could freely enter boxes to interact via a screen. In the social condition, baboons played in separate adjacent boxes without being able to see the conspecific's choices. First, by fitting the probability to get a good answer and the one to switch targets with GLMs, we found that the difference in hierarchical status between subjects significantly explained choices. Higher dominance level (determined by ELO-score) led to switching less between targets. As the algorithm encouraged stability, this also led to a better performance. The marginal effects showed that in the social condition, the switching frequency of baboons decreased faster than in the non-social condition, and they performed better in the social condition. A comparison between classical learning models (without social dominance) then showed differences between the social and non-social conditions: the most frequent strategy in the non-social condition was a simple heuristic. In the social condition, a mentalizing influence model fitted better observed behavior. Thus, despite the same algorithm running in both conditions, baboons mentalized their partners when interacting with a conspecific, while they apply simpler heuristics when interacting with their environment. Finally, to check how social rank affected decision learning, we modified the previous mentalizing model to account for the differences in dominance levels, and found that it indeed fitted better behavioural data. More specifically, this model suggests that the difference in dominance levels modulate the learning rate of their mentalized partner:

dominant subjects mentalize slow learning partners. This therefore opens new perspectives on the computational mechanisms behind the influence of social hierarchy on coordination learning.

O5.2 Lesions of Nucleus Accumbens Shell abolish Socially Transmitted Food Preferences

Irina Noguer Calabus[1], Sandra Schäble[1], Tobias Kalenscher[1]

[1]Heinrich-Heine-Universität Düsseldorf

Objective: As a social species, an important faculty for us is social learning, i.e., the acquisition of adaptative behaviors by observing or interacting with a conspecific. In rodents, the Socially Transmitted Food Preference (STFP) task is a well-established paradigm to study social learning, consisting of overwriting the endogenous food preference by the preference of the conspecific. Many brain regions underly the acquisition of the STFP paradigm, but the role of the Nucleus accumbens shell (NAcS) in this task remains unknown. Since the NAcS is an important area in the reward system for adaptative responses, we hypothesized that the NAcS plays a crucial role in the STFP paradigm. Methodology: 36 actor rats received either bilateral NAcS lesion or sham surgery. In the STFP task, two food options were provided to actor rats and their consumption patterns were measured to determine their individual food preference. Then, demonstrator rats were fed with the respective non-preferred food option and interacted with an actor rat for 20 minutes. After the interaction, the actors had, again, access to both food options and their consumption pattern was measured anew. Two control tasks were performed to rule out potential confounds: anxiety was measured in an open field test, and an Odor Discrimination task was carried out to control for odor recognition deficits. A Choice index (CI) was computed before and after the interaction to quantify the normalized food preference for one option over the other. The CIs were analyzed using a general linear model (GLM) with two groups (NAcS lesion and Sham lesion) as between-subjects factor and Time (pre and post interaction) as within-subjects factor. The histological analysis is still ongoing at the time of abstract submission. Results: The GLM revealed a significant main effect of Time (Pre vs. Post; $F(1, 34) = 7.706$, $p = 0.009$) on the CI as well as a significant Time x Lesion (Sham vs. Lesion) interaction effect ($F(1, 34) = 6.969$, $p = 0.012$), suggesting that the contact with the demonstrator rat affected actor rats' food choices differentially for each experimental group: while choices of the originally non-preferred food increased after demonstrator contact in the Sham group, this change in choice was not observed in the Lesion group. These results were not explained by anxiety levels or a deficit in odor recognition. Conclusion: This experiment suggests that NAcS lesions result in a deficit in socially transmitted reward revaluation. These results provide new information about the areas underpinning social learning and contribute to the understanding of social influence on choice and preference.

O5.3 Biased algorithms produce biased humans: the consequences of human-AI collaboration

Moshe Glickman[1], Tali Sharot[1]

[1]University College London (UCL)

Objective: Artificial Intelligence (AI) bias is the underlying prejudice in the data used to develop algorithms, resulting in discrimination. For example, a hiring algorithm used by Amazon was inevitably trained to favor men over women and an algorithm used in US hospitals to estimate the need for additional medical care favored white patients over black. However, the gravity of AI bias may have been underestimated, as critical decisions are often made jointly by AI and humans. Such collaboration

introduces an opportunity for AI bias to negatively impact human assessment. Here, we examine whether collaborating with a biased AI makes human judgement more biased. Methods: Over 3 experiments participants (N1=30, N2=90, N3=90) collaborated with different algorithms. In E1 the collaborative task was motion discrimination. In E2 & E3 it was assessing the valence expressed by a group of faces. We built 3 simple algorithms which provided either a systematically biased response (biased AI), an accurate response (accurate AI) or a noisy response (noisy AI). On each trial participants first indicated their response and then observed the algorithm's response. The participant then decided how much weight to assign to their own response and to the AI's response in the final joint decision. Each participant played one block with each algorithm. Before being exposed to the algorithms participants played the tasks on their own, allowing us to quantify baseline ability. Results: Participant own judgements became biased when playing with the 'biased AI' (E1: $t(29)=3.01$, $p<.01$, E2: $t(89)=2.93$, $p<.01$, E3: $t(89)=3.65$, $p < .001$). The more participants collaborated with the 'biased AI', the more they were influenced by its bias (E1: $t(2606)=2.41$, $p<.05$, E2: $t(7826)=3.93$, $p<.001$, E3: $t(8174)=4.89$, $p<.001$). Participants, however, were unaware that they became biased (all ps>.4). A computational model revealed that playing with the biased AI biased the weights subjects gave to the available pieces of evidence (E2: $t(2720)=5.78$, $p<.001$, E3: $t(2960)=5.01$, $p<.001$) and altered the response bias (E1: $t(867)=8.67$, $p<.001$). Finally, using simulations we characterize a path by which bias could spread among populations via AI-human interactions. Conclusions: Several algorithms providing input into decisions related to health care, finance and the legal system, have shown to be biased. Here, we show that collaborating with a biased AI causes subjects to become biased themselves, an effect subjects are unaware of. This effect could trigger a feedback loop (i.e., biased data generating biased AI, which generates biased data) resulting in inaccurate judgements and discriminatory actions.

O5.4 Separable neurocognitive changes underlie the development of communicative reasoning in adolescence

Cong Wang[1], Menghan Cong[1], Qingtian Mi[1], Guihua Yu[1], Min Hong[1], Skylar Hong[1], Yanjie Su[1], Lusha Zhu[1]

[1]Peking University

Objective Humans possess a remarkable ability to read between the lines. Prior research suggest that decoding what is meant from what is said involves several intertwining cognitive operations, such as inhibiting the literal interpretation of communicative signals, and inferring the speaker's most probable intention possibly through simulating the speaker's goal-directed choices over candidate signals to best deliver an intention in context. Despite recent progress in understanding the neural bases of this important class of social function, little is known about when and how humans acquire this ability, and whether basic cognitive functions such as executive control and social reasoning would interact with one another in support of communicative development. Combining model-based fMRI with cross-sectional and longitudinal data, we examined the neurobehavioral changes in communicative reasoning during the transition in and out of adolescence, a critical period for the development of both cognitive and social functioning. Methods A total of 316 subjects (ages 8-24) participated in the experiment, among which 53 fMRI subjects were scanned twice, spanning 1.7-year. Communicative ability was assessed using a well-established referential game and a computational model previously used to elucidate the neurocomputational mechanism of communicative reasoning in adults. Results The accuracy of communicative reasoning continuously increased with age in both cross-sectional ($r = 0.55$, $P < 0.001$) and longitudinal data (T2-T1 = 8.29%, $P < 0.05$). The age-related improvement is related to (i) an

adolescent-specific development in the egocentric processing of contexts ($F(3,137) = 5.67$, $P < 0.005$), subserved by changes in sensory cortical activity (small-volume PFWE < 0.05); and (ii) an adolescent-nonspecific development in reasoning speaker's intention-action contingency ($r = 0.49$, $P < 0.001$), subserved by changes in the ventromedial prefrontal cortex engagement (cluster-wise PFWE < 0.05). Importantly, the relative contribution of these processes is modulated by changes in the involvement of frontoparietal network (FPN), at both cross-sectional and longitudinal levels. Conclusions These results suggest a prolonged development of communicative reasoning over the entire adolescence period. It likely involves a delicate balance between the sensory processing and social reasoning, which is correlated with the development of FPN functioning. These results offer a potential mechanism by which the maturation of communicative reasoning is shaped by separable, yet interacting neurocognitive processes.

Poster Spotlight Abstracts

Poster Spotlight #1

PS1.1 (Poster #1-G-6) From GWAS to Molecular and Social Pathways to Subjective Well-Being: An Empirical Analysis of the Canadian Longitudinal Study on Aging

Mike Cisneros-Franco[1], Catherine Paquet[2], Daiva Nielsen[1], Alain Dagher[1], Patricia Silveira[1], Laurette Dube[1]

[1]McGill University, [2]Universite Laval

Adaptive behavior in the face of a changing environment is essential for achieving and maintaining well-being throughout life. Genome-wide association studies (GWAS) have been used to generate polygenic risk scores (PRS) for subjective well-being, based on numerous loci linked to physical/mental health outcomes (e.g., predictors of metabolic syndrome and affective disorders). However, this approach does not consider the molecular pathways through which dopamine genes related to decision-making and emotional regulation are expressed in specific brain regions, or how the networks they form with other genes may interact with the social environment in impacting subjective well-being in older adults. We examined gene-by-social environment (G*E) interactions in subjective well-being (CESD-10 scale) in the Canadian Longitudinal Study of Aging (age 45-75 years; n = 13950, 6733 women). In addition to a GWAS-derived PRS for subjective well-being, we used a novel expression-based polygenic risk score reflecting variations in the expression of the DRD4 gene network in striatum (ePRS-DRD4). Enrichment analysis of the ePRS was performed to evaluate biological pathways. Using latent profile analysis, participants were classified based on their social environment (network size, social cohesion, social support, social participation, and objective social isolation). This approach, based solely on the social environment measures irrespective of outcomes, identified three distinct profiles: low-, medium-, and high-social network support (18/40/42%, respectively), with the low- and medium-support groups exhibiting overall lower levels of subjective well-being relative to the high-support group. Significant G*E interactions were found, such that high PRSs (related to higher propensity for poor subjective well-being) in the low-support group were associated with worse subjective well-being. In addition, lower striatal expression of the DRD4 gene network as measured by ePRS-DRD4 was associated with lower subjective well-being in the low-support profile. Genes in the Striatum-ePRS-DRD4 network mainly overlapped with other gene

sets related to cardiometabolic health (blood pressure, body-mass index, type 2 diabetes), suggesting common pathways for physical and mental well-being. Differences in genetic background, including the striatal DRD4 gene co-expression network, may moderate social environmental influences on older adults' well-being. Further population-based studies on G*E interactions may provide valuable insights in identifying the most susceptible older adults and developing age-appropriate evidence-based strategies to strengthen their social support networks.

PS1.2 (Poster #1-D-7) A gut feeling: how your gut and brain determine your choices

Aline Dantas[1], Elisabeth Bruggen[1], Peiran Jiao[1], Alexander Sack[1], Teresa Schuhmann[1]

[1]Maastricht University

Recent research has shown that gut microbiota can influence the interaction between the central and the enteric nervous system via the gut-brain axis (GBA). Neural activity in brain regions linked to basic emotional as well as cognitive processes have already been shown to be affected by manipulations of the gut microbiota. Whether such gut microbiota manipulations also affects human decision making, however, remains largely unknown. Based on previous studies in animal models and indications from neuroimaging studies, we here test the effects of the gut-brain axis on decision-making under risk and intertemporal choices. In a placebo-controlled double-blinded design, with two sessions separated by 28 days, during which participants received daily doses of probiotics (or placebo), we investigated whether the prolonged and controlled intake of probiotics affects risk-taking behavior and intertemporal choices using incentivized games. We found a significant decrease in risk-taking behavior and increase in future-oriented choices in the probiotics group, as compared to placebo. These findings provide first direct experimental evidence for the functional role of the microbiota-gut-brain axis on decision-making, creating a path for potential clinical applications and allowing for a better understanding of the underlying neural mechanisms of risk-taking behavior and intertemporal choices.

PS1.3 (Poster #1-F-8) Neural representations of others' traits predict social decisions

Kenji Kobayashi[1], Joseph Kable[1], Ming Hsu[2], Adrianna Jenkins[1]

[1]University of Pennsylvania, [2]University of California, Berkeley

OBJECTIVE: People perceive others along core trait dimensions, such as warmth and competence, and these perceptions are known to influence decisions about how to treat and interact with members of different social groups. However, much remains unknown about how human brains represent others' traits in the service of social decision making. Using fMRI and representational similarity analysis (RSA), we examined neural representations of others' traits and their contribution to decision making in an economic game. METHODS: Participants (n = 32) decided how to share money with members of twenty different social groups (e.g., "Nurse"; "Irish") in a Dictator Game. On each trial, they were first presented with a recipient in one of the social groups and then with two monetary allocation options, one with an equal allocation and the other with an unequal allocation. The unequal option allocated more money to the participant than the recipient on some trials (advantageous inequity) and more money to the recipient than the participant on other trials (disadvantageous inequity). RESULTS: At the behavioral level, the recipient's perceived warmth and competence had dissociable effects on choices in the Dictator Game. Participants were more averse to advantageous inequity when the recipient's warmth was higher, irrespective of their competence; conversely, they were more averse to disadvantageous

inequity when the recipient's competence was higher, irrespective of their warmth. At the neural level, a searchlight RSA revealed that the receiver's perceived traits were represented in a two-dimensional space of warmth and competence in the right temporoparietal junction (TPJ) and the left superior temporal sulcus (STS), regions long associated with mentalizing, and the left orbitofrontal cortex (OFC), a region associated with inference-based goal-directed behavior. Critically, an additional RSA revealed that activation patterns in the OFC were predictive of individual decisions; individual-level representational dissimilarity in the OFC was correlated with subsequent monetary allocation choices in the advantageous and disadvantageous inequity contexts. CONCLUSIONS: The OFC is widely known to contribute to goal-directed decisions by representing latent features of the environment or task that are critical for outcome inference, but the literature has largely focused on non-social decisions. Our results suggest that the representation of others' traits in the OFC influence monetary allocation decisions in a context-dependent manner, and more generally, that the OFC plays critical roles in inference or imagination of decision outcomes in the social domain.

PS1.4 (1-I-9) Information-seeking during COVID-19

Nathan Torunsky[1], Kara Kedrick[1], Iris Vilares[1]

[1]University of Minnesota - Twin Cities

Objectives: In the context of a global pandemic, understanding the motivations behind information-seeking can provide useful insights for public health education and intervention. Sharot & Sunstein (2020) recently proposed that information-seeking is driven by estimates of the hedonic, instrumental, and cognitive utility of potential resources. Additionally, past research has shown that information-seeking can be driven by emotional states and tendencies. The goal of the present study was to identify predictors of COVID-19 information-seeking and examine the relationship between individual differences in emotion and the calculation of information value. Methods: We recruited 187 undergraduate participants from the University of Minnesota during the fall of 2020. Participants completed five rounds of an information-seeking task in which they read about a topic related to COVID-19, rated their expected hedonic, instrumental, and cognitive utility of an article excerpt on that topic, and decided whether to seek or avoid the excerpt. After finishing the information-seeking task, participants completed self-report questionnaires assessing their general mood state (i.e. depression, anxiety, and stress), intolerance of uncertainty (IU), engagement in COVID-19-preventative behavior (e.g. mask-wearing), and trust in science and scientists. In addition to Spearman correlations, we fit mixed-effects logistic regression models to predict participants' information seeking behavior, and used BIC and cross-validation to select the best-fitting model. Results: Model selection yielded a model with only a single predictor of information-seeking: cognitive utility ($\beta = 1.2$ [0.90, 1.6], $p = 2.2e\text{-}11$, OR = 3.5 [2.5, 5.1]). Depression, anxiety, and IU were not correlated with overall information-seeking or participants' estimates of information utility, but preventative behaviors positively correlated with information-seeking ($r = .21$, $p = .006$). Trust in science and scientists was positively correlated with information-seeking, but failed to reach significance, possibly due to poor variation in scores ($r = .14$, $p = .064$). Trust in science and scientists was strongly correlated with engagement in COVID-19 preventative behaviors ($r = .31$, $p = 3.3e\text{-}5$). Conclusions: In an undergraduate sample, information-seeking about COVID-19 appears to be primarily driven by the expected cognitive utility yield of a particular resource rather than by hedonic or instrumental utility or individual differences in emotional states. Additionally,

self-reported preventative behaviors appear to correlate both with information-seeking and trust in science and scientists.

PS1.5 (Poster #1-I-10) Value Certainty in Drift-Diffusion Models of Preferential Choice

Douglas Lee[1], Marius Usher[2]

[1]Institute of Cognitive Sciences and Technologies, National Research Council of Italy, [2]Tel Aviv University

The drift-diffusion model (DDM) is widely used and broadly accepted for its ability to account for binary choices (in both the perceptual and preferential domains) and for their response times (RT), as a function of the stimulus or the option values. The DDM is built on an evidence accumulation to bound concept, where, in the value domain, a decision maker repeatedly samples the mental representations of the values of the options until satisfied that there is enough evidence in favor of one option over the other. As the signals that drive the evidence are derived from value estimates that are not known with certainty, repeated sequential samples are necessary to average out noise. The classic DDM does not allow for different options to have different levels of variability in their value representations. However, recent studies have shown that decision makers often report levels of certainty regarding value estimates that vary across choice options. There is therefore a need to extend the DDM to include an option-specific value certainty component. We present several such DDM extensions and validate them against empirical data from four previous studies. The data support best a DDM version in which the drift of the accumulation is based on a sort of signal-to-noise ratio of value for each option (rather than a mere accumulation of samples from the corresponding value distributions). This DDM variant accounts for the positive impact of value certainty on choice consistency and for the negative impact of value certainty on decision time, both of which are present in the empirical data.

Poster Spotlight #2

PS2.1 (Poster #2-E-36) Deliberative evaluation in intertemporal choice is shaped by experiment structure

Nidhi Banavar[1], Aaron Bornstein[1]

[1]University of California - Irvine

Objective: Modeling choice behavior or response time in randomized, or unstructured, intertemporal choice (ITC) experiments often assumes an absence of trial-level sequential dependencies. A natural extension of popular theories of ITC that involve simulating the future (Peters & Buchel, 2010) imply that already computed future values could be cached and re-used, especially if an individual makes similar choices in sequence (Dasgupta et al, 2018). Here we test how sensitive our inferences about the components of deliberative evaluation are to such "spillover" effects through response time modeling. Methods: Subjects (n = 482) participated in a text-based ITC task with no explicit sequential dependencies (Hunter, Bornstein, Hartley, 2018). Each subject made a sequence of 102 incentive-compatible decisions where they chose between a smaller amount of money today or a larger amount of money in the future. To link response time modeling with stimulus properties, we developed a Bayesian implementation of a Drift Diffusion Model approximation (Bogacz et al, 2006). We then applied

a hierarchical framework that allows for distinguishing whether response times are influenced by spillover between ostensibly unrelated trials. This framework can simultaneously discern whether there is meaningful evidence for spillover on multiple different aspects of the cognitive processes under examination - e.g. on discount factor, drift rate, or bias. Results: Nearly one in three subjects show meaningful evidence (BF > 3) of sequential effects in the drift rate and/or overall bias towards patience or impulsivity. These effects are driven by both cross-trial value and delay differences, as well as their interaction. Approximately half of these effects manifest on an individual's bias parameter, to the degree that parameter interpretation changes: sequential effect adjusted parameters show individuals inferred to be patient (prefer Later option) in the standard approach as actually having an overall preference for the immediate option. Our model also successfully captures a wide range of individual differences in both the presence and absence of spillover in a non-sequential task. Conclusions: Our results highlight the importance of explicitly accounting for the fact that humans process information sequentially. We demonstrate that when behavioral data is analyzed with this fact in mind, the estimated variables in higher order cognition can change not only in magnitude but also in interpretation. This is critical for improving the mapping between parametric estimates and real-world behavior.

PS2.2 (Poster #2-I-37) Escaping Arrow's Impossibility by Interpersonal Comparison of Neural Utility

Kaosu Matsumori[1], Kazuki Iijima[1], Yukihito Yomogida[1], Kenji Matsumoto[1]

[1]Tamagawa University

Aggregating welfare across individuals to reach collective decisions is one of the most fundamental problems in our society. Interpersonal comparison of utility is pivotal and inevitable for welfare aggregation, because if each person's utility is not interpersonally comparable, there is no rational aggregation procedure that simultaneously satisfies even some very mild conditions for validity (Arrow's impossibility theorem). However, scientific methods for interpersonal comparison of utility have thus far not been available. Here, we have developed a method for interpersonal comparison of utility based on brain signals, by measuring the neural activity of participants performing gambling tasks. We found that activity in the medial frontal region was correlated with changes in expected utility, and that, for the same amount of money, the activity evoked was larger for participants with lower household incomes than for those with higher household incomes. Furthermore, we found that the ratio of neural signals from lower-income participants to those of higher-income participants coincided with estimates of their psychological pleasure by ?impartial spectators?, i.e. disinterested third-party participants satisfying specific conditions. Finally, we derived a decision rule based on aggregated welfare from our experimental data, and confirmed that it was applicable to a distribution problem. These findings suggest that our proposed method for interpersonal comparison of utility enables scientifically reasonable welfare aggregation by escaping from Arrow's impossibility and has implications for the fair distribution of economic goods. Our method can be further applied for evidence-based policy making in nations that use cost-benefit analyses or optimal taxation theory for policy evaluation.

PS2.3 (Poster #2-G-38) Does endogenous variation in stress modulate risk and time preferences?

Evgeniya Lukinova[1], Jeffrey Erlich[1]

[1]NYU Shanghai

Objective: It has been argued that one dimension of the cycle of poverty is that poverty is a state of chronic stress. Poor decisions, made under chronic stress, might include carrying high-interest loans, failure to buy health insurance, gambling or drug use. As such, these decisions can contribute to the cycle of poverty. More specifically, a few studies suggest that increased stress may lead to more risk aversion and steeper delay discounting. While the deleterious effects of chronic stress on brain function are well established, much less is known about how chronic stress influences financial decision making. Methods: Here, in a longitudinal design within six weeks period we aimed to incorporate biological mechanisms to improve our understanding of how stress influences economic decisions. We used a combination of decision-making tasks, stress questionnaires, saliva and hair samples within-subject (N=41). We assessed time and risk preferences using hierarchical Bayesian techniques to both pool data and allow heterogeneity in decision making and compared those to cortisol levels and self-reported stress. Results: We found only weak links between endogenous variation in stress and model-based estimates of risk and time preferences. In particular, we found that stress was correlated more strongly with risk preferences, rather than time preferences: cortisol levels in the month preceding the risk task correlated with increased risk tolerance. Also, we found model-free task measures in the short delay task to be moderately related to both hair and saliva cortisol, as well as the stressful life events questionnaire measure. For example, we observed that the stress level one month before the task or the life change units were negatively correlated with the proportion of later choices. Finally, similarly to previous results about subjects' earnings, we established that the curvilinear relationship was preferred to the linear one for waiting time: when cortisol level increased slightly, people waited longer in seconds but when stress increased to higher levels, people waited less. Conclusions: Our results do not directly support the hypotheses that increased stress leads to steeper discounting and increased risk aversion. Instead, we found that endogenous subclinical variation in cortisol only weakly correlates with economic preferences. Moreover, the strongest link we found was that increased cortisol levels correlated with decreased risk aversion, the opposite of our prediction from the literature. Additional multi-task and multi-stress studies with higher N are required to resolve the discrepancies between these results and existing literature.

PS2.4 (Poster #2-I-39) Malleability of human altruism across choice contexts due to social cues

Lisa Bas[1], John-Dennis Parsons[1], Anita Tusche[1]

[1]Queen's University

Introduction. The willingness to act altruistically varies across contexts. What mechanisms drive decisions to bear substantial costs to benefit others in some settings but not in others? Here, we combine an altruistic choice task with formal modeling approaches to examine the role of social cues across contexts. We hypothesize that cues of others' emotional distress - signaling need - may increase generosity. We also investigate differences in people's sensitivity to social factors. Methods. Subjects (N=58) performed an online version of the dictator game. On every trial, subjects chose between two offers that affected their own and another person's payoff. A generous choice was defined by accepting the offer that benefitted the other at a cost to oneself or by rejecting the offer benefitting oneself at a cost to another. Subjects indicated their preference using a computer mouse. Continuous measures of the cursor position offer a window into the moment-by-moment construction and temporal dynamics of the choice process. To examine the role of social contextual factors, subjects performed the task under two conditions: 'distress' (2 s exposure to dynamic facial expressions of sadness before each trial) or

'control' (no social facial cues) (2 blocks each, 132 trials). Subjects also completed an emotion recognition task and self-report measures of trait differences in mentalizing. Results. Exposure to others' distress increased the generosity compared to the control condition. However, we observed substantial individual differences in the impact of social cues on choices. Subjects with a tendency to spontaneously take others' point of view (better mentalizers) were more affected by social cues during altruistic choice. They were also naturally slower at integrating other's benefits into their decision, possibly due to sufficient leeway for change in the processing speed of others' gains across contexts. Subjects who were responsive to social cues in the dictator game also reported higher real-world donations. This finding suggests that differential sensitivities to social factors may extend to real-world contexts. Notably, emotion recognition performance (d') was comparable for more or less responsive individuals. Thus, our results are unlikely due to differential abilities to correctly identify others' emotional states. Conclusion. The results provide insights into the mechanisms driving variance in altruism across contexts. Our work also points to characteristics of the decision-maker that modify the impact of contextual factors on social behavior. Together, our findings can help to advance research on the malleability of altruism.

PS2.5 (Poster #2-E-40) Investigating the link between neural reward reactivity and attention

Nitisha Desai[1], Allison Londerée[1], Eunbin Kim[1], Dylan Wagner[1], Ian Krajbich[1], Kentaro Fujita[1]

[1]The Ohio State University

Self-control failures--the undermining of long-term goals for short-term rewards--are implicated in many societal problems, including obesity, substance abuse, and poor financial decisions. Previous research has shown that self-control failure is related to greater activity in the brain's reward regions (e.g., ventral striatum) when exposed to temptations. By contrast, greater activity in the regions involved in inhibitory control (e.g., inferior frontal gyrus) supports successful resistance to temptation. Along with these neural mechanisms, visual attention has also been shown to play a role in self-control decisions. Research has found that attentional biases towards short-term rewards are associated with self-control failures. Here--using a multi-modal approach that combines functional MRI (fMRI), eye-tracking, and diffusion modeling--we investigate the link between neural reward reactivity to appetitive stimuli and attentional biases to short-term rewards. In our experiment, 26 dieters participated in an fMRI scan followed by an eye-tracked choice task. Dieting status was confirmed by commonly used metrics of dieting concern (i.e., Revised Restraint Scale). We used the scanning session to measure individual differences in reward cue-related activity to appetitive foods. After the fMRI session, we assessed subjects' self-control by having them make binary choices between healthy and unhealthy foods. There were 300 trials in total; 100 "conflict" trials directly pitted a healthy food against an unhealthy food. We found that activity in the left nucleus accumbens (NAcc) predicted the tendency to look at the unhealthy item first ($r(24)=0.55$, $p=0.003$), whereas activity in the inferior frontal gyrus (IFG) was related to the probability of looking at the healthy item first ($r(24)=0.39$, $p=0.047$). To further explore the relationship between reward reactivity and attention, we fit an attentional drift-diffusion model. Modeling results revealed that activity in the right ventromedial prefrontal cortex activity was negatively correlated with starting point bias (z) towards the healthy option ($r(24)=-0.41$, $p=0.038$) and with the attentional discount (θ) on the not-looked-at option ($r(24)=-0.46$, $p=0.017$). Our results suggest greater NAcc activity may be associated with a bias to initially seek out short-term rewards, while greater IFG activity may be associated with the opposite, a bias to seek out the options consistent with long-term goals. These novel

findings provide a first step towards identifying a mechanism--early attentional biases--by which neural reward reactivity impairs self-control and neural inhibitory activity promotes it.

Poster Spotlight #3

PS3.1 (Poster #3-E-68) Goal-dependent recalibration of hippocampal representations facilitates self-control

Micah Edelson[1], Todd Hare[1]

[1]University of Zurich

Objective: The neurobiological interactions between mnemonic and decision-making processes are important for determining individual behavior. One domain where this may be especially relevant are decisions that require self-control because they rely on the ability to select and control the internal information utilized when making a choice. For example, when dieting and maintaining a goal to eat healthfully, successful self-controllers may be able to prioritize the recall of a subset of experiences in which chocolate cake was less tasty than usual or a salad was particularly delicious. Theoretically, this would result in mnemonic representations of the past and/or future forecasted tastiness that are more similar for chocolate cake and salad (because the tastiness of salad is increased and that of chocolate cake decreased). We refer to this information control process as goal-dependent recalibration and predict that it should occur in the hippocampus given that region's well-established role in episodic memory. Methods: We tested our hypothesis by conducting representational similarity analyses of fMRI activity within the hippocampus while humans (N = 44) performed blocks of 50 food rating or dietary choice trials. Results: We found that people do flexibly adjust their multivariate hippocampal representations of the palatability of food items in a goal-dependent fashion. On average, the entire sample of participants showed significant goal-dependent recalibration of taste representations in the hippocampus between trials in which the goal was to rate the tastiness versus rate the healthiness of a food item (Cohen's d= 0.69, posterior probability = 0.999). Critically, this goal-dependent recalibration was also associated with actual choices in the face of self-control challenges. In fact, participants with high, relative to low, levels of self-control did represent palatable and unpalatable food items more similarly within the hippocampus when facing self-control challenges (Cohen's d = 0.8, posterior probability = 0.989) and the amount of goal dependent recalibration was strongly predictive of self-control success (out-of-sample accuracy = 69%, probability > chance = 0.996). Conclusions: Our results provide new insights into the processes involved in self-control. They indicate the need to incorporate memory and hippocampal representations into mechanistic frameworks of self-control alongside the longstanding focus on i) regulatory functions supported by the prefrontal cortex and, ii) antagonistic interactions between prefrontal cortex and subcortical brain regions.

PS3.2 (Poster #3-E-69) Does COVID-related stress affect self-control and the ability to make healthy food choices?

Marie Falkenstein[2], Felix Nitsch[3], Leonie Koban[2], Aiqing Ling[4], Tobias Kalenscher[3], Hilke Plassmann[1]

[1]INSEAD, [2]Sorbonne University, [3]Heinrich-Heine-Universität Düsseldorf, [4]University College Dublin

Objective: Experimentally inducing acute stress has been previously found to impair self-control. Research on the effects of chronic stress on self-control is sparse but research on more general decision-making suggests a negative impact. The body responds to stressors by stimulating the hypothalamic-pituitary-adrenal (HPA) axis releasing cortisol which is commonly known as the stress hormone. The ratio of cortisol to cortisone can be used to assess the activity of the HPA axis. Here, we investigated whether the Covid-19 crisis is an exogenous stressor that impacts self-control. Methods: Our study consisted of two parts: (1)An online study carried out across China and France summer 2020. Participants were asked to make choices between foods and report their chronic stress levels, those since the pandemic and at the moment amongst other items. (2)Participants send a hair strain to extract their hair cortisol and cortisone levels to compute their ratio. Data of N=553 (36% CN, 78% females, Mage= 26, SD=6.8) were analyzed. A food choice task was used in which participants rated healthiness and tastiness of different food items. Thereby, food choice sets were built for each participant with half of the trials showing foods where the tastier item was less healthy and thus self-control was required to make healthy choices. Results: To assess whether the COVID-19 pandemic led to increased stress, we compared the self-reported chronic stress-levels with prior literature using the same scale (TICS). We found that the mean reported TICS in our study was significantly higher than the score in the previous literature (MStudy=16.55, SDStudy=5.96 vs. Mliterature=9.4, SDliterature=6.3, $t(3019)=24.1$, $p<.001$). Further, we tested whether the ability to make healthier choices when self-control is required is linked to stress. We find a very low, albeit significant, correlation with pandemic-induced stress ($r(539) =-.09$, $p =.04$). When looking at the concentration of cortisol in relation to cortisone in hair, we also find a small negative marginally significant correlation with the percentage of healthy choices ($r(469)= -.09$, $p= .055$). Conclusion: We found that the higher chronic COVID-19-induced perceived stress, the lower people's ability to make healthy food choices when self-control is required. We found the same trend when using a biological marker of chronic stress, i.e., hair cortisol concentration. However, these negative correlations were very small and did not hold when investigating both countries separately. Together, our results showed much smaller effect sizes as compared to controlled lab studies investigating the impact of acute stress on self-control.

PS3.3 (Poster #3-I-70) Parallel representation of context and multiple context-dependent values in ventro-medial prefrontal cortex

Nir Moneta[1], Mona Garvert[2], Hauke Heekeren[3], Nicolas Schuck[1]

[1]Max Planck Institute for Human Development Berlin, [2]Max Planck Institute for Human Cognitive and Brain Sciences, [3]Freie Universität Berlin

Objective: The ventromedial prefrontal-cortex (vmPFC/mOFC) is known to signal values that guide decisions. But how preferable available options are depends on one's current task. Goal-directed behavior, which involves changing between different task-contexts, therefore requires knowing how valuable the same options will be in different contexts. Whether such a multiplicity of values exists and how it might interact with value-free information about task structure in the same area is unknown. We tested whether multiple task-dependent values influence behavior and asked if they are integrated into a single value representation or are co-represented in parallel within vmPFC signals. Methods: Thirty-five participants alternated between tasks in which stimulus color or motion predicted rewards, while undergoing fMRI. On each trial, participants were cued to focus only on one context and had to choose the cloud leading to the largest reward. This meant that the same set of visual features was associated

with different values in different task states. Although in each trial task states were associated with different values, across trials they had the same values and were dissociated from perceptual input. Results: Behavioral analysis indicated that participants did not merely perform a value-based choice among currently relevant values. Rather, both reaction times and accuracy showed they retrieved the values of irrelevant features and computed the resulting counterfactual choice. Using multivariate fMRI classifiers we investigated the probability distributions over values, which shed light on the complexity of representations in vmPFC. We found that the vmPFC maintains representations of the value associated with the current context/task state and the hypothetical value that would be expected in the alternative task state. Current task context could also be decoded from the same region. Crucially, we show these representations interact: expected and hypothetical values compete, and stronger task state representations are associated with stronger representations of its value. Finally, our results suggest a link between neural representations of task states and their associated values to their influence on behavior. Conclusions: We shed new light on vmPFC's role in decision making and bridge between previous perspectives of its role: first, it is involved in mapping observations onto a mental map, i.e. inferring the current task state; second, it is involved in computing the value expectations that result from the inferred state and third it also computes the hypothetical value that would result if one would have been in a different task state.

PS3.4 (Poster #3-J-71) Browsing under threat: high-level features of web searches altered during the pandemic and predicted population stress levels

Christopher Kelly[1], Bastien Blain[1], Tali Sharot[1]

[1]University College London

Objective: The 2020 pandemic generated a set of practical & mental challenges. To overcome these, people turned to the internet. Here, we examined how the high-level features of information people sought online (e.g. instrumental utility of information, valence) changed during the pandemic. To do so, we analyzed over a trillion internet searches submitted weekly before & during the pandemic in the UK & US. We related these measures to (i) stress levels reported weekly by 20K individuals & (ii) population-level mobility data extracted from mobile phones. We theorized that during the pandemic people will be more likely to seek answers to questions that can guide action & increase comprehension of the world around them. Methods: We calculated the % of "How to" & "How do" Google searches submitted in the UK & US every week from 2017 to 2021. These searches likely result in information that can guide action (e.g. "how do you install Zoom?"). We did the same for "What" & "Why" searches. These likely increase comprehension (e.g. "what is Zoom?"). We also calculated the %positive minus %negative words entered for the most popular searches by matching words to an emotion lexicon. We then examined (i) how these features changed during the pandemic, & (ii) how these changes related to population stress levels reported weekly by 20K people in the UK. We dissociated the effects of stress from the effect of being confined to one's residence by controlling for mobility rates quantified from mobile phone data. Results: There was a sharp increase in proportion of searches that could guide action (i.e. "How to/How do") during the pandemic relative to before ($p = 0.0001$), which was positively associated with weekly population-stress levels, controlling for mobility rates ($p = 0.0001$). This suggests that under stress people seek information with high instrumental utility. This feature of web searches predicted population stress levels better than searches for specific stress-related terms (e.g. "anxiety"; $p = 0.003$). Mobility rates, but not stress levels, were associated with the proportion of "What/Why" searches,

suggesting that home confinement may have increased the desire for general knowledge. Searches were also more negative. Results were replicated in UK & US. Conclusion: The increase in proportion of searches that can guide action may have contributed to human resilience in response to the pandemic. Markedly, this feature of web searches predicted population stress levels better than searches for specific terms (e.g. "psychiatrist"). An intriguing possibility is that the described method could be used for monitoring population stress levels beyond the pandemic.

PS3.5 (Poster #3-H-67) Episodic decision-making via a process of cascading episodic sampling (CASES)

Achiel Fenneman[1], Alan Sanfey[2]

[1]Radboud University Nijmegen / Rhine-Waal University of Applied Science, [2]Radboud University Nijmegen

Daily life often presents us with complex decisions, the properties of which typically prevents us from relying on previous statistically-learned outcome values. Instead, we have to rely on our episodic memories to determine the actions that are available to us - and which of these is most attractive choice. In order to model this memory-based decision-making, we need to answer two related but independent questions: which memories do we rely upon, and how do we use these memories to make a decision. Though previous models and theories of memory-based decision-making exist, none provide a psychologically plausible answer to both of these questions. Here we describe such a model based on recent findings in memory and decision-neuroscience. This model is built around a bottom-up cascading spread of memory activation: the activation of a memory results in a simultaneous feed-forward and feed-back process. The feed-forward process generates evidence to either positively or negatively bias the action previously associated with this memory. If this process does not result in sufficient evidence to warrant a decision, then a feed-back pattern-completion process reinstates the features that make up the activated memories. These reinstated features then activate an additional set of memories, trigging the next feed-forward and feed-back processes, until a decision has been reached. We formalize this process in a model of Cascading Episodic Sampling (CASES), and show that this non-random sampling process is robust to large number of stored memories, and provides a faster decision speed as the number of memories increases. This model parsimoniously captures, and expands upon, a range of key findings and previous models in the topic of memory-based decision-making.

Poster Abstracts

Poster Session #1

1-C-11 Trial-Level Changes in Reward Positivity/Feedback-Related Negativity During a Competitive Game Using Multilevel Modeling

Brian Howatt[1], Matthew Wisniewski[1], Michael Young[1]
[1]Kansas State University

Objective: In the present study we examined trial-by-trial changes in the magnitude of the reward positivity/feedback-related negativity (RewP/FRN) event-related potential (ERP) component during a

competitive game using multilevel modeling. While feedback-locked ERP components are typically derived by averaging voltage amplitudes across trials to obtain a single magnitude estimate, recent research has shed light on the theoretical and empirical benefits of modeling how the magnitude changes at the trial-level. Here, we share empirical results on how individuals' RewP/FRN component magnitudes change as a function of their competitor's predictability, and discuss implications on current theories of feedback processing. Methods: Participants (N=57) played Rock, Paper, Scissors against computerized opponents programmed using a Q-Learning reinforcement learning algorithm. To simulate different human-like opponent strategies, algorithms were randomly assigned model parameters (a learning rate, alpha, and an inverse temperature, tau) that produced random choices (alpha = tau = 0), moderately predictable choices (alpha = 0.3, tau = 5) or highly predictable choices (alpha = 1.0, tau = 10) -- by 'predictable' choices, we mean the tendency for the model to engage in a win-stay, lose-shift strategy. The opponent predictability was a between-subjects variable (random n = 20; moderately predictable n = 20; highly predictable n = 17) and each participant played only one opponent. All feedback-locked ERP analyses were conducted using multilevel linear regressions with the model intercept, trial slope and outcome slope (which represents the RewP/FRN magnitude) varying across participants as random effects. Results: The magnitude of the RewP/FRN component (200-300 ms) at electrode location FCz decreased for the highly predictable and random groups across trials, but slightly increased for the moderately predictable group. Further analyses suggested that the reduction in RewP/FRN magnitude for the highly predictable group was likely contaminated by the feedback-related positivity (Pe) and P3 components. However, the magnitude changes for the moderately predictable and random groups did not appear to be contaminated by the Pe and P3 components. Conclusions: Overall, our study further affirms that the RewP/FRN magnitude systematically changes across trials and indicates that such changes depend on a competitor's predictability. We consider our findings within the context of different theories of feedback processing and learning, such as reward prediction error and salience prediction error.

1-D-7 A gut feeling: how your gut and brain determine your choices

Aline Dantas[1], Elisabeth Bruggen[1], Peiran Jiao[1], Alexander Sack[1], Teresa Schuhmann[1]
[1]Maastricht University

Recent research has shown that gut microbiota can influence the interaction between the central and the enteric nervous system via the gut-brain axis (GBA). Neural activity in brain regions linked to basic emotional as well as cognitive processes have already been shown to be affected by manipulations of the gut microbiota. Whether such gut microbiota manipulations also affects human decision making, however, remains largely unknown. Based on previous studies in animal models and indications from neuroimaging studies, we here test the effects of the gut-brain axis on decision-making under risk and intertemporal choices. In a placebo-controlled double-blinded design, with two sessions separated by 28 days, during which participants received daily doses of probiotics (or placebo), we investigated whether the prolonged and controlled intake of probiotics affects risk-taking behavior and intertemporal choices using incentivized games. We found a significant decrease in risk-taking behavior and increase in future-oriented choices in the probiotics group, as compared to placebo. These findings provide first direct experimental evidence for the functional role of the microbiota-gut-brain axis on decision-making, creating a path for potential clinical applications and allowing for a better understanding of the underlying neural mechanisms of risk-taking behavior and intertemporal choices.

1-D-13 Surprising sounds influence decision making

Gloria Feng[1], Robb Rutledge[1]
[1]Yale University

Adaptive behavior requires rapid flexible responses to surprising sensory events. However, although common in daily life, such events are often behaviorally irrelevant. It is unknown whether surprising sensory events elicit consistent behavioral responses, even when sensory events are task irrelevant and behavioral responses are potentially maladaptive. In four experiments (n=800), we tested whether task-irrelevant auditory surprises influence how humans choose between risky and safe options. Using sequences of tones ending in an auditory oddball, we found that surprising sounds preceding option presentation increased risk taking irrespective of whether the risky options featured potential gains or losses. Moreover, sensory surprise increased switching away from the option chosen on the previous trial. These findings were consistent across two studies (each n=200), and were unaffected by whether the location of risky options changed. We manipulated different aspects of auditory surprise in two additional studies (each n=200) and found that these effects were dissociable, suggesting different underlying mechanisms. Lower-order surprise, related to stimulus novelty, increased risk taking. Higher-order surprise, based on learning which auditory sequences were more common, increased option switching. Using computational modeling, we show that both effects are explained by changes in uncertainty bonus and perseveration parameters, which do not depend on option values. Our findings align with previous studies characterizing sensory prediction errors related to dopamine, a neuromodulator also strongly implicated in decision making. More broadly, our findings demonstrate that surprising sounds common in the natural world may influence many of the decisions we make in daily life.

1-D-14 Emotional Influence on Information Processing in Decisions under Uncertainty

Silvia Lopez-Guzman[1], Santiago Sautua[2]
[1]National Institute of Mental Health NIMH, [2]Universidad del Rosario

Introduction Decision makers undoubtably make use of information about historical outcomes of their choice options to make decisions. For example, a financial investor usually observes previous returns of several financial assets before choosing a portfolio; and a physician learns how several medical treatments have affected a given patient in the past before choosing again a combination of treatments for the patient. Importantly, the way in which this information may be processed and incorporated into a decision could be affected by emotional state. While emotional states are well known to impact decision-making, less is understood about how they influence decisions when information about prior outcomes is present. Methods In two laboratory experiments (N = 443), we investigate the effect of incidental emotions on decision-making under ambiguity, both in the presence and absence of information about past outcomes. In the experiments, the subjects make investment decisions and report their expectations of success after watching a short video clip with either negative or positive valence content. Results Incidental emotions induced by the video clip have differential effects on expectations and investments depending on whether the subjects have information about prior outcomes. When subjects lack such information, a negative emotional state reduces investments (β=-0.082 $P<0.05$), consistent with prior evidence that some negative emotions reduce risk-taking. By contrast, following strong bad news about past outcomes, a negative emotional state attenuates the tendency to reduce expectations (just bad news β=-0.169 $P<0.01$, bad news + negative emotion β=0.061

P<0.05) and investments (just bad news β=-0.96 P<0.01, bad news + negative emotion β=0.002 P=N.S) among "conventional decision makers." In addition, in this situation, a significant minority of subjects who experience negative emotion engage in "contrarian behavior," increasing expectations and investments despite the bad news. Conclusions Overall, our findings show that new information about historical outcomes could eliminate or even reverse the effects that emotions exert on decision-making under static uncertainty. These results underscore the need for additional research investigating how the interaction between emotional states and new information affects decision-making under uncertainty.

1-D-15 Incentives and arousal modulate the perception of value in risky choice

Abdelaziz Alsharawy[1], Sheryl Ball[2], Alec Smith[2]
[1]Princeton University, [2]Virginia Tech

Objective: The brain must use limited computational resources to make decisions. The principle of efficient coding implies that value representations are context dependent. In particular, more frequently encountered payoffs are perceived more accurately. We hypothesized that these representations would be influenced by incentives (real vs. hypothetical) and by affective states (arousal and valence). Our objective was to test how incentives and emotions modulate value perception in risky choice. Methods: Participants (N=70) completed a series of 600 decisions choosing between two options: 1) a lottery with a 50% chance of a positive payoff and a 50% of a zero payoff and 2) a sure payoff. The payoffs were sampled from a distribution with a narrow range (low volatility - LV- condition) in half of the trials and from a distribution with the same mean but a wider range (high volatility- HV- condition) in the other half of the trials, with the order of the conditions counterbalanced across participants. Crucially, participants were assigned to realize real payment, based on a randomly determined decision, from either the LV or HV conditions. Participants were informed that payoffs in the other volatility condition were instead hypothetical. During the session, participants were trained to classify emotions and regularly reported their emotional experience (arousal/valence). Results: We confirm the efficient coding hypothesis only for participants assigned to receive real payment from the LV condition, where perception was more sensitive to changes in payoffs under the LV condition compared to the HV one. On the other hand, participants assigned to receive real payment from the HV condition displayed comparable sensitivities to changes in payoffs across volatility conditions. Moreover, we find that self-reports of arousal, and not valence, were significantly higher in real payment condition. We then compute the mean difference in arousal for each participant across volatility conditions. We find that sensitivity to changes in the risky option's payoffs did not decline for participants experiencing amplified arousal levels in the HV condition. In addition, we find that both reaction time and risk aversion were higher in the real payment condition and were strongly linked to individual differences in arousal. Conclusion: We find that stronger (real) incentives modulate the perception of value and increase self-reports of arousal. The efficient coding hypothesis seems to hold best under weak incentives and low levels of arousal. Our results demonstrate the importance of incentives and emotional experiences in the adaptation of perceptual processing of value.

1-D-16 Structural coherence of white-matter fiber tracts converging on the nucleus accumbens is associated with different aspects of risk preference

Loreen Tisdall[1], Kelly MacNiven[2], Josiah Leong[3], Renato Frey[1], Jörg Rieskamp[1], Ralph Hertwig[4], Brian Knutson[2], Rui Mata[1]

[1]University of Basel, [2]Stanford University, [3]University of Arkansas, [4]Max Planck Institute for Human Development

Functional responses in the Nucleus Accumbens (NAcc) to risk- and reward-related cues can predict real-life risk-taking behavior. Since NAcc activity is driven by neurotransmission from connected brain regions, the structural coherence of tracts projecting to the NAcc may also predict risk preference. In a sample of 125 healthy human adults (54% female, 25.19±2.6 years old), we identified projections from Medial Prefrontal Cortex (MPFC), Anterior Insula (AIns), Amygdala (Amy), and Ventral Tegmental Area (VTA) to the NAcc, in order to assess whether their coherence was associated with risk preference. To quantify risk preference, we used latent variables that were previously derived in a large, independent study on the psychometric structure of risk preference. To precisely identify and characterize the tracts, we combined probabilistic tractography in individuals' native space with constrained spherical deconvolution. Statistical analyses targeted a commonly used index of tract coherence, fractional anisotropy (FA), and included permutation-based testing methods to derive summary FA coherence metrics for each tract in each individual. Based on previous work, we predicted and preregistered a negative association for a general risk preference factor, R, with right AIns-NAcc (H1) and bilateral MPFC-NAcc tract coherence (H2). For a latent variable capturing impulsivity, we predicted a negative association with bilateral VTA-NAcc tract coherence (H3), and a positive association with bilateral Amy-NAcc tract coherence (H4). We further predicted an association between R and bilateral Amy-NAcc tract coherence (H5), but tested a bidirectional hypothesis due to the heterogeneity of previous findings. As predicted, VTA-NAcc tract coherence was negatively associated with a psychometric factor capturing impulsivity (FA b=-2.82, SE=1.12, t=-2.53, p=0.013), and Amy-NAcc tract coherence was positively associated with this factor (b=3.06, SE=1.12, t=-2.53, p=0.013). Further, MPFC-NAcc tract coherence was positively associated with R, although a negative association was predicted (b=2.92, SE=1.26, t=2.31, p=0.023). Although an association was predicted, AIns-NAcc and Amy-NAcc tract coherence were not associated with R. All analyses controlled for effects of age and gender. Taken together, our results replicate and extend previous results by demonstrating that the coherence of projections to the NAcc are differentially associated with aspects of risk preference. Although functional mechanisms remain to be tested, structural projections to the NAcc reveal an exciting map for future research on the physiological underpinnings of risk preference.

1-D-31 Understanding How Individual Beliefs About Risk Contribute to Ambiguity Aversion

Jeffrey Dennison[1], David Smith[1]
[1]Temple University

Objective: We often need to make decisions without knowing how likely certain outcomes are, even though we may have some information about similar situations. For example, we may not know how likely we are to get a good coffee from a coffee shop we have never been to, but we have visited other coffee shops and maybe even other coffee shops from the same chain. Such decisions--which are made without risks being described--highlight the concept of ambiguity. Although much of the prior work in economics has not considered the role of individual beliefs about risk in decisions involving ambiguity (i.e., maxmin expected utility), recent work has begun to examine the role of individual beliefs about risk (i.e., recursive utility). In this project, we test predictions of a recursive utility model under conditions of ambiguity and compare it to predictions from the maxmin expected utility model. Methods: We create a naturalistic set of compound lotteries, in order to manipulate individual beliefs about the risks in

ambiguous decisions, and combine it with a manipulation of ambiguity used for maxmin utility models. This combination allows for tests that better compare the two models than previous explorations. Participants (N=27) make decisions about three differently colored lotteries under conditions where the probabilities are displayed (risk) and when those probabilities are fully or partially hidden (ambiguity). While making decisions during the risk condition, each color is associated with a different second-order distribution of probabilities. During the ambiguous condition, only the color of the lottery and the range of possible risk (i.e., 0%-100% or 25%-75%) are shown. Participants also completed questionnaires assessing mood symptoms and substance use. Results: We predicted that lottery colors associated with more variable second-order probabilities during risky decision would be devalued during the ambiguous decision. We used a hierarchical logistic regression and found that when the range of risk was fully extended individuals devalued lotteries with greater second-order variance, t(23)=-2.5, p=0.013. However, the range of risk had no effect for lotteries with narrow second-order distributions t(23)=-0.66, p=0.5. Conclusions: Ambiguity aversion has been associated with clinically important outcomes for addiction (Konova et al., 2020). Our experiment suggests that decisions under ambiguity rely on the combination of tolerance to second-order risks and our beliefs about them. Previous work has treated these two constructs as one parameter. Dissociating these parameters would be fruitful for future experiments or interventions targeting each const

1-D-32 Modelling reversed framing effects in experience: Reference-based overweighting of extreme outcomes as a learning mechanism

Y.L. Doug Dong[1], Kevin da Silva Castanheira[1], Christopher Madan[1], Ross Otto[1]
[1]McGill University

BACKGROUND: The description-experience gap depicts that the framing effect, which is typically observed in one-shot decisions from description, is reversed in repeated and feedback-based risky decisions. This apparent reversal in risk-preferences is difficult to account for under prospect theory, an influential and ubiquitous explanation for the framing effect. A growing body of work suggests that decision-makers use contextual information as reference points to evaluate outcomes as gains or losses. For example, an objectively small loss can be reframed as a subjective gain relative to a context of large losses. At the same time, one influential account - the extreme-outcome rule - predicts that the reversed framing effect arises as the result of overweighting extreme outcomes in memory; however, this theory has yet to be formalized computationally. Our goal is to develop a computational account for the reversed framing effect observed in decisions from experience. METHODS: We propose a series of models that combine prospect theory's reference-dependent overweighting, and mechanisms of value updating from reinforcement learning (RL) algorithms. These models extend traditional RL models by learning the expected value of a given context, updated with every outcome experienced trial-by-trial through Q-learning. With this contextual information as a reference point, these models then non-linearly weight the trial outcome to produce subjective values as in prospect theory. Critically, the more extreme an outcome is (i.e., more deviant from the context reference), the more it is overweighed. In the end the outcome's subjective value is used to update the expected value of the chosen stimulus through Q-learning. To examine model performances, we use the pre-existing data from a well-characterized choice task that documents the reversed framing effect. We simulate the proposed models, standard RL algorithms, and their variations with parameter values reported in the literature and evaluate each model's individual fit to participant data (N=133). RESULTS: We show that our novel models both predict the reversed framing effect in simulation and outperform standard models (e.g.,

delta-rule RL models) in terms of goodness-of-fit, even after taking model complexity into account. CONCLUSION: Together we formalize the extreme-outcome rule as a computational model and suggest that overweighting outcomes relative to a learned reference point can predict the reversed framing effect commonly observed in decisions from experience.

1-D-34 Imprecise Probabilistic Inference from Sequential Data

Arthur Prat-Carrabin[1], Michael Woodford[1]
[1]Columbia University

How people form expectations about possible future outcomes is a fundamental question in economics and cognitive neuroscience. Following the emergence of the "rational expectations" hypothesis, it is commonly assumed in economics that beliefs are optimally updated, through Bayesian inference. This assumption is also a common benchmark in the cognitive sciences. However, under-reaction to new information ("conservatism") is found in some inference tasks, while at the same time over-reaction is also sometimes reported. To investigate human departures from optimal Bayesian inference, we design an inference task in which subjects are asked to estimate the probability of a binary event occurring, on the basis of successive realizations of the event. Our experimental design allows us to examine the bias, variability, and autocorrelation in the estimates of subjects who are repeatedly presented with the same short sequences of observations. We find under-reaction of estimates to the presented evidence after only a few observations, and at the same time over-reaction to the evidence after a longer sequence of observations. We test and reject several models proposed in the literature on human inference (e.g., that subjective probabilities are non-linear transformations of the ones implied by the evidence; or that people correctly update, but starting from a prior different than the one that is correct for the experimental environment). In addition, the autocorrelation in estimates suggests that the noise in a subject's response largely results from the imprecision of their mental representation of the decision situation, rather than arising only at the stage of response selection. Finally, we find that subjects' estimates do not conform to key properties of the estimates of any Bayesian observer, even one with an imprecise memory. Overall, subjects' responses are consistent with a "noisy counting" model of probability estimation, in which they maintain an imprecise count of the net excess of one outcome over the other, but keep little track of the total amount of evidence accumulated. This model yields a better fit (as measured by the Bayes Information Criterion) than other inference models found in the literature, and it reproduces the behavioral patterns found in data, including the under-reaction to short series of observations and the over-reaction to longer series, and the autocorrelation in responses. The noisy-counting model allows subjects to give relatively reasonable responses to the task while economizing greatly on both the attention that they must pay to their current situation and the degree of control that they exert over their precise response.

1-E-17 Temporal Variations in Caloric Intake and Body Weight Among Daily Users of a Mobile Self-Monitoring Application: Relationship between Tracking Compliance and Long-Term Weight Loss

Katherine Labonté[1], Bärbel Knäuper[1], Laurette Dubé[1], Nathan Yang[2], Daiva Nielsen[1]
[1]McGill University, [2]Cornell University

Objective: Data from calorie tracking applications (apps) enable investigation of the role of self-monitoring on adherence to a caloric budget and weight loss over time. We evaluated the effect of season and body mass index (BMI) on adherence to caloric intake goals using data from one calendar

year (2016) from n=9372 individuals (50% male) who tracked their caloric intake with an app daily. We also examined the relationship between tracking of body weight in the app and weight loss. Methods: Linear mixed models were used to assess the effects of season (winter, spring, summer, fall) and BMI group at baseline (normal, overweight, obese) on the mean number of calories exceeding users' daily budget (kcal/d). Analyses were stratified by gender (F: female; M: male) due to physiological differences in body weight, and were adjusted for age and reliability of caloric intake (reliable if ≥800 kcal/d). Associations between weight self-monitoring and change in body weight were analyzed using a subset of users (n=5808) who entered their weight at least once per month, once per season, or once in both winter and fall. To correct for multiple testing, $p<0.01$ was used for statistical significance. Results: Mean caloric intake was lowest in winter for all users, but only users with obesity adhered to their budget during this season. Caloric intake increased from winter to spring for all users (mean kcal/d over budget: +23.7 (F), +39.7 kcal (M), $p<0.001$), remained stable for the summer, and further increased among users with obesity in the fall (+14.0 kcal (F), +11.8 kcal (M), $p\leq0.005$). Users with obesity had a smaller deviation from budget overall (mean difference: +40.0 (F), +54.5 kcal (M)) than those with overweight (+119.3 (F), +158.8 kcal (M)) and normal weight (+128.5 (F), +126.3 kcal (M)) ($p<0.001$). Users who entered their body weight less than at least once per season had no significant change in weight at the end of the year (F: -0.2, M: +0.8 lbs, $p\geq0.017$). Among those with more regular entries, weight loss was observed for users with obesity (F: -11.2, M: -9.5 lbs, $p<0.001$) and overweight (F: -3.2, M: -1.2 lbs, $p\leq0.001$), but not for users with normal weight (F: +0.3, M: +1.1 lbs, $p\geq0.018$). Conclusions: Among highly consistent app users, only users with obesity adhered to a caloric budget early in a calendar year. Caloric intake was lowest at this time for all users, and observed temporal variations may reflect depletion of self-regulatory resources. Regular self-monitoring of body weight with calorie tracking may be more effective for weight loss. Differences in app use motive could explain the lack of weight loss for users with normal weight.

1-F-18 Social learning from advice and observation: how trust and reputation affect what we share and how we learn from others

Uri Hertz[1]
[1]*University of Haifa*

Socially-transmitted information is the basis for our species' extraordinary ecological success. Social information can be gleaned via two main routes: either by observing how others behave or by following explicit advice or recommendations. Advice taking is different from observational learning, as it relies on evaluation of the advisors' intentions and trustworthiness, beyond evaluation of competence. This differentiation is important, as it has clear implications on the way we understand social information transmission in humans, beyond animal models which rely mostly on observational learning, and the way psychiatric disorders and social conditions where trust is undermined affect social learning. Nevertheless, it is not clear to what extent people treat these social information sources differently, nor how this might vary with the underlying tendency to trust other people. I will present two studies which used a learning task and reinforcement learning models to investigate the way people learn from observation and advice, and the factors affecting advice giving. In the first study, participants learned which of two options was more rewarding. On trial 5, participants either received advice from, or observed the choice of, an expert. Participants were more likely to immediately follow advice than to copy an observed choice but this was dependent upon trust in the adviser: highly paranoid participants were less likely to follow advice in the short-term. Reinforcement learning modelling revealed two

distinct patterns regarding the long-term effects of social information: some individuals relied fully on social information and chose the suggested option in all remaining choices, whereas others reverted to trial-and-error learning. Advice was privileged in the long term effect as well. In a second study, participants followed a similar learning task, and had the opportunity to broadcast their decisions as advice to other players on every trial. We found that participants were overall calibrated in their advice, and were more willing to share their information later in learning, and when the task was easier. This finding supports the privileged status of advice over observation. However, some participants were more calibrated than others, and likelihood of advice was dependent on self-esteem, social anxiety and gender. Taken together, our findings highlight the way social infrastructure like reputation, and individual characteristics like trust and self-perception, may affect prevalence and fidelity of socially-transmitted information.

1-F-8 Neural representations of others' traits predict social decisions

Kenji Kobayashi[1], Joseph Kable[1], Ming Hsu[2], Adrianna Jenkins[1]
[1]University of Pennsylvania, [2]University of California, Berkeley

OBJECTIVE: People perceive others along core trait dimensions, such as warmth and competence, and these perceptions are known to influence decisions about how to treat and interact with members of different social groups. However, much remains unknown about how human brains represent others' traits in the service of social decision making. Using fMRI and representational similarity analysis (RSA), we examined neural representations of others' traits and their contribution to decision making in an economic game. METHODS: Participants (n = 32) decided how to share money with members of twenty different social groups (e.g., "Nurse"; "Irish") in a Dictator Game. On each trial, they were first presented with a recipient in one of the social groups and then with two monetary allocation options, one with an equal allocation and the other with an unequal allocation. The unequal option allocated more money to the participant than the recipient on some trials (advantageous inequity) and more money to the recipient than the participant on other trials (disadvantageous inequity). RESULTS: At the behavioral level, the recipient's perceived warmth and competence had dissociable effects on choices in the Dictator Game. Participants were more averse to advantageous inequity when the recipient's warmth was higher, irrespective of their competence; conversely, they were more averse to disadvantageous inequity when the recipient's competence was higher, irrespective of their warmth. At the neural level, a searchlight RSA revealed that the receiver's perceived traits were represented in a two-dimensional space of warmth and competence in the right temporoparietal junction (TPJ) and the left superior temporal sulcus (STS), regions long associated with mentalizing, and the left orbitofrontal cortex (OFC), a region associated with inference-based goal-directed behavior. Critically, an additional RSA revealed that activation patterns in the OFC were predictive of individual decisions; individual-level representational dissimilarity in the OFC was correlated with subsequent monetary allocation choices in the advantageous and disadvantageous inequity contexts. CONCLUSIONS: The OFC is widely known to contribute to goal-directed decisions by representing latent features of the environment or task that are critical for outcome inference, but the literature has largely focused on non-social decisions. Our results suggest that the representation of others' traits in the OFC influence monetary allocation decisions in a context-dependent manner, and more generally, that the OFC plays critical roles in inference or imagination of decision outcomes in the social domain.

1-G-6 From GWAS to Molecular and Social Pathways to Subjective Well-Being: An Empirical Analysis of the Canadian Longitudinal Study on Aging

Miguel Cisneros-Franco[1], Catherine Paquet[2], Daiva Nielsen[1], Alain Dagher[1], Patricia Silveira[1], Laurette Dube[1]
[1]McGill University, [2]Universite Laval

Adaptive behavior in the face of a changing environment is essential for achieving and maintaining well-being throughout life. Genome-wide association studies (GWAS) have been used to generate polygenic risk scores (PRS) for subjective well-being, based on numerous loci linked to physical/mental health outcomes (e.g., predictors of metabolic syndrome and affective disorders). However, this approach does not consider the molecular pathways through which dopamine genes related to decision-making and emotional regulation are expressed in specific brain regions, or how the networks they form with other genes may interact with the social environment in impacting subjective well-being in older adults. We examined gene-by-social environment (G*E) interactions in subjective well-being (CESD-10 scale) in the Canadian Longitudinal Study of Aging (age 45-75 years; n = 13950, 6733 women). In addition to a GWAS-derived PRS for subjective well-being, we used a novel expression-based polygenic risk score reflecting variations in the expression of the DRD4 gene network in striatum (ePRS-DRD4). Enrichment analysis of the ePRS was performed to evaluate biological pathways. Using latent profile analysis, participants were classified based on their social environment (network size, social cohesion, social support, social participation, and objective social isolation). This approach, based solely on the social environment measures irrespective of outcomes, identified three distinct profiles: low-, medium-, and high-social network support (18/40/42%, respectively), with the low- and medium-support groups exhibiting overall lower levels of subjective well-being relative to the high-support group. Significant G*E interactions were found, such that high PRSs (related to higher propensity for poor subjective well-being) in the low-support group were associated with worse subjective well-being. In addition, lower striatal expression of the DRD4 gene network as measured by ePRS-DRD4 was associated with lower subjective well-being in the low-support profile. Genes in the Striatum-ePRS-DRD4 network mainly overlapped with other gene sets related to cardiometabolic health (blood pressure, body-mass index, type 2 diabetes), suggesting common pathways for physical and mental well-being. Differences in genetic background, including the striatal DRD4 gene co-expression network, may moderate social environmental influences on older adults' well-being. Further population-based studies on G*E interactions may provide valuable insights in identifying the most susceptible older adults and developing age-appropriate evidence-based strategies to strengthen their social support networks.

1-H-19 A micro-genesis account of longer-form reinforcement learning (gain-calmness, loss-restlessness) in structured and unstructured environments

Ben Dyson[1], Ahad Asad[1]
[1]University of Alberta

Fundamental reinforcement learning principles such as win-stay and lose-shift represent outcome-action associations between consecutive trials (trial n-1 and n). Longer-form expressions of the tendency to continually repeat previous actions following positive outcomes, and the tendency to continually change previous actions following negative outcomes, have been identified as win-calmness and lose-restlessness, respectively. Across 10 experiments (n = 356) using the simple game of Rock, Paper, Scissors, we tested a micro-genesis account of these phenomena by examining sequential contingencies

across trial n-2, n-1 and n using simple game spaces. At a group level, we found no evidence of win-calmness and lose-restlessness when wins could not be maximized (unexploitable opponent). Similarly, we found no evidence of win-calmness and lose-restlessness when the threat of win minimization was presented (exploiting opponent). In contrast, we found evidence of win-calmness (but not lose-restlessness) when win maximization was made possible (exploitable opponent). At a participant level, we confirm that individual win rates determined the degree of win-calmness and lose-restlessness only in contexts were win rates could be maximized. The data identify the mechanisms that allow for the development of longer-form reinforcement learning principles and demonstrate the relative flexibility in decision-space afforded by positive outcomes, and the relative inflexibility in decision-space following negative outcomes.

1-H-33 Anxiety increases the emotional impact of negative prediction errors during learning

Chang-Hao Kao[1], Olivia Jorasch[1], Robb Rutledge[1]
[1]Yale University

Learning is driven by prediction errors that capture the difference between beliefs and experience (Sutton & Barto, 1998). Anxiety reduces the ability to appropriately adapt to changing environments (Browning et al., 2015), suggesting that anxiety may lead to inappropriate responses to prediction errors. Here we ask whether anxiety influences the behavioral or emotional impact of prediction errors in simple learning tasks. In two experiments (n=200 and n=199), participants chose between options with different but high probabilities of reward (0.9 and 0.6) in multiple blocks of trials and indicated how happy they were periodically during the task. Average performance in both experiments was high (68% and 66%) and performance was not correlated with anxiety assessed using the Generalized Anxiety Disorder-7 questionnaire (both $p>0.5$). We fitted participants' choices using a reinforcement learning model with separate learning rates for positive and negative prediction errors (mean pseudo-r2 of 0.34 and 0.32). Learning rates for both positive and negative prediction errors were uncorrelated with anxiety (all $p>0.1$). Trial-by-trial value estimates and prediction errors were derived from this learning model. We then fitted participants' happiness ratings, and estimated the influence on happiness of value estimates and prediction errors for previous trials (mean r2 of 0.54 and 0.54), consistent with previous research showing that the recent history of expectations and prediction errors influences happiness in risk taking tasks that do not involve learning (Rutledge et al., 2017). The weight for positive prediction errors in both experiments was not significantly correlated to anxiety (Spearman rho=0.08, p=0.25; rho=0.14, p=0.06). The weight for negative prediction errors in both experiments was significantly correlated to anxiety (rho=0.19, p=0.006; rho=0.20, p=0.004). Our results suggest that anxiety may influence the emotional impact of prediction errors even in stable environments where anxiety does not influence behavior.

1-I-10 Value Certainty in Drift-Diffusion Models of Preferential Choice

Douglas Lee[1], Marius Usher[2]
[1]Institute of Cognitive Sciences and Technologies, National Research Council of Italy, [2]Tel Aviv University

The drift-diffusion model (DDM) is widely used and broadly accepted for its ability to account for binary choices (in both the perceptual and preferential domains) and for their response times (RT), as a function of the stimulus or the option values. The DDM is built on an evidence accumulation to bound concept, where, in the value domain, a decision maker repeatedly samples the mental representations

of the values of the options until satisfied that there is enough evidence in favor of one option over the other. As the signals that drive the evidence are derived from value estimates that are not known with certainty, repeated sequential samples are necessary to average out noise. The classic DDM does not allow for different options to have different levels of variability in their value representations. However, recent studies have shown that decision makers often report levels of certainty regarding value estimates that vary across choice options. There is therefore a need to extend the DDM to include an option-specific value certainty component. We present several such DDM extensions and validate them against empirical data from four previous studies. The data support best a DDM version in which the drift of the accumulation is based on a sort of signal-to-noise ratio of value for each option (rather than a mere accumulation of samples from the corresponding value distributions). This DDM variant accounts for the positive impact of value certainty on choice consistency and for the negative impact of value certainty on decision time, both of which are present in the empirical data.

1-I-20 Context-dependent choice and evaluation in real-world consumer behavior

Ross Otto[1], Sean Devine[1], Aaron Bornstein[2], Kenway Louie[3]
[1]McGill University, [2]University of California - Irvine, [3]NYU

A body of work spanning neuroscience, economics, and psychology indicates that decision-making is context-dependent, which means that the value of an option depends not only on the option in question, but also on the other options in the choice set--or the 'context'. While context effects have been observed primarily in small-scale laboratory studies with tightly constrained, artificially constructed choice sets, it remains to be determined whether these context effects take hold in real-world choice problems, where choice sets are large and decisions driven by rich histories of direct experience. Here, we investigate whether valuations are context-dependent in real-world choice by analyzing a massive restaurant rating dataset (Yelp.com; 4.2 million ratings). We find that users make fewer ratings-maximizing choices in choice sets with higher-rated options--a hallmark of context-dependent choice-- and that post-choice restaurant ratings also varied systematically with the ratings of unchosen restaurants. These context effects are well characterized, computationally, by a Divisive Normalization model of choice, in which subjective expectations of options' quality are computed in a context-dependent manner. Furthermore, in a follow-up laboratory experiment using hypothetical choice sets derived from the real-world choice sets, we find that individuals' expectancy ratings are also contextually modulated in accordance with a Divisive Normalization view, with the consequence that, all else being equal, more valuable contexts systematically decreased expectancy ratings down, providing corroborating evidence for a general mechanistic-level account of these effects. Taken together, we find compelling evidence for context-dependent valuation in real-world choice settings, manifesting both in choice and subjective ratings of options.

1-I-21 Real-world decision making: does it require an ecological-social construct of self?

Gina Kemp, Georg Northoff[1], Laurette Dubé[2]
[1]University of Ottawa Institute of Mental Health Research (IMHR), [2]McGill University

Current models of decision-making have limited fidelity in predicting real-world behavior, especially relevant to subjective preferences. One key dynamic feature in decision making is the self and its identity, as recently proposed by Berkman's Self-Identity Model. In this model Berkman et al. stipulate that during a situation of choice, alignment with self-identity positively augments the value of certain

options. While this model integrates the self into the valuation process as self-related content, it does not account for other facets of the self (i.e., its cultural and social-relational nature). This is important because integrating the self provides an opportunity to interface biological life with social, cultural, and ecological life in the interaction with the environment. We present a model which integrates different aspects of the self which allows for the empirical investigation of real-world decision making as a complex process. Specifically, we will present a nested model operating across and integrating different timescales based on the scale-free nature of the self. In this model, cognitive-perceptual processes, which include valuation, are nested within an agentic aspect of the self called the minimal self. This aspect is characterized by an immediate "here and now" sense of body, movement, and thought ownership with the following characterizations: 1) First, its temporal nature allows for its extension across time in a way that can be captured empirically in evolutionary biology, genetics, and computer science. 2) Second, it is content-less and is in the form of temporal dynamics of a set of regions in the brain's midline, related to the Default-Mode Network (DMN). 3) Lastly, the minimal-self can be extended beyond the boundaries of particular agent-like persons to its social, cultural, and ecological contexts. Such contexts are characterized by their own temporal dynamics exhibiting fluctuations in timescales that partially overlap with those of the minimal self - one can thus speak of a "common currency" of self, brain, and environment (Northoff et al. 2020). Drawing on and extending the psychological concept of Gestalt theory, one can say that the self, i.e., minimal self, can be conceived as a figure against and amidst the socio-cultural environmental context as its background. Together, the here proposed concept of the minimal self allows shedding a novel light on the complexity of our decision making. Due to its temporally-dynamic and multi-scale nature, such an extended concept of the minimal self is well equipped to take into view the social-cultural, biological, and ultimately evolutionary basis of our decision making.

1-I-22 Eye-tracking in the assessment of subjective experience measures in virtual reality environment

Michal Gabay[1], Tom Schonberg[1]
[1]Tel-Aviv University

Introduction: The implementation of virtual reality (VR) has been spreading in many research fields and healthcare applications. In order to better personalize the presented content, there is need to identify the subjective experience at each time point. Previous studies with eye-tracking (ET) have demonstrated its links to preferences (e.g. Shimojo et al., 2003, Krajbich et al., 2010), therefore collecting ET data during VR experiences could be useful for passive identification of value and for the purpose of content personalization. Objective: We aimed to develop an objective and quantitative method for the identification of subjective preference measures based on ET data in VR. Methods: We used VR goggles adapted with a binocular high-speed ET device (Tobii-installed VR HTC-Vive) in a dynamic VR task we developed in Unity©. Participants explored a scene and were asked to search and look at objects of two types: pets and control shapes that appeared in pre-defined locations in random order. During the VR task, ET as well as behavioural data were recorded, and features were extracted in order to detect the rankings of each stimulus preference, valence and arousal. Results: Based on a pilot sample with n=23 participants we hypothesised that the median distance of gaze (Med_Dist) while looking at an object will be linked to preference and valence rankings on a larger second sample of 53 participants. Hypotheses and sample size were pre-registered prior to data collection. We analyzed the data using mixed linear models accounting for object size, as well as the gazed time in order to control for a possible mere exposure effect (Zajonc, R. B. & Markus, H, 1982). We found a significant interaction of Med_Dist and

type of objects (Pets/Shapes) in both preference (t[87.4]=-5.8, p=2.84e-7) and valence (t[140.9]=-5.7, p=2e-7). In addition, in line with our hypotheses a significant positive relation was found between Med_Dist and preference (t[49.14]=4.7, p=2.3e-5) and valence (t[66.4]=6.2, p=2.6e-8) rankings of pets. Tests were Bonferroni corrected for 3 ranking models. The gazed time had no effect on rankings, thus eliminating the alternative explanation that mere exposure drove increased rankings (p>0.2 in all models). Conclusions: We developed a novel objective passive and quantitative method for the identification of subjective experience measures of preference and valence while viewing single items, based on a VR built-in eye-tracker. Our results suggest that ET can be used as a potential biomarker for detecting individual preference and pleasantness while engaged in VR, and in the future allow real time updating of VR content for various applications.

1-I-23 A novel smartphone app to induce and study habits in humans

Rani Gera[1], Segev Barak[1], Tom Schonberg[1]
[1]Tel Aviv University

Objective: Habits are a fundamental feature of behavior, affecting almost every aspect of everyday life. In contrast to animal research, habits in humans are relatively poorly characterized, mainly as there are no well-established procedures to induce habits in humans. We rationalized that the typical characteristics of laboratory timeframe constrains, artificial context and structured dense training sessions may not be optimal to study habit formation. Methods: We developed a novel smartphone app, free of the traditional laboratory constraints, yet well-controlled and informed by basic scientific principles. The app is implemented as a game in which participants enter the app freely and voluntarily 24/7 to earn a reward (outcome). The ?cover? story is that they need to dig gold on a remote planet. In an initial sample (n=32) we randomly assigned participants to three groups: a short-training group (5 days; n=12) and two long-training groups (12 days; long training groups data were collapsed for analysis purposes; n=20). On the 4th or 11th day of the short and long training groups, respectively, the outcome is devalued (the storage gets full until the next day) to test participants? sensitivity to outcome devaluation, an accepted criterion to distinguish goal-directed form habitual behavior. Entries to the app despite outcome devaluation are considered habitual responses. Results: Our preliminary results indicate that our novel smartphone app captures the individual variation in habit expression (Mean habitual entries across all participants=5.31, SEM=1.9). Importantly, comparing changes in response rate following outcome devaluation between the groups has shown that participants were more likely to respond habitually following long training compared with short training (Mann-Whitney U=77, p=0.048). Furthermore, the chances of committing at least one habitual response were substantially larger in the longer training groups compared to short training (p=0.073, Fisher's exact test). Conclusions: We established a new smartphone app. Based on the pilot sample it seems suitable for measuring habit acquisition and expression at the group and individual levels. Our method may provide a promising way to represent the related behavioral and cognitive processes in more naturalistic conditions.

1-I-24 Speed-accuracy tradeoffs in a disinhibition-based neural circuit model of decision-making

Bo Shen[1], Paul Glimcher[1], Kenway Louie[2]
[1]New York University Grossman School of Medicine, [2]New York University

Objective: Speed-accuracy tradeoffs (SAT) in decision-making have long been a phenomenon of interest in psychology and neuroscience. As a signature of decision dynamics, the SAT has been considered a

necessary benchmark for neural decision circuits but how such dynamics arise is unknown. Sequential models of decision-making predict behavioral SAT on the assumption of the shifting of decision criteria. However, this assumption is contradicted by recent neurophysiological findings on neural decision thresholds and does not capture the empirical evidence for baseline elevation or gain tuning in neural dynamics. Here we examine how well a novel biologically-based circuit - the local disinhibition decision model (LDDM) - captures the behavioral and neural signatures of SAT across different datasets. Methods: The LDDM implements recurrent excitation and global gain-control inhibition drawn from existing models and incorporates a biologically plausible element of local disinhibition that governs "top-down" control of the network. The model was implemented as a dynamical system capturing firing rate responses. We tested this model in capturing the behavioral and neural dynamics in a range of empirical datasets: 1) behavior and neural dynamics of SAT in monkeys (Heitz and Schall, 2012; Hanks et al., 2014; Thura & Cisek, 2016, 2017), 2) SAT under individual differences in working memory (Heitz and Engle, 2007), and 3) causal manipulation of NMDA antagonism on the behavior of SAT (Shen et al., 2010). Results: We show that the LDDM captures the behavioral tradeoff between reaction time (RT) and choice accuracy by controlling the intensity of local disinhibition during the process of decision making. The best-fitting models to the behavior capture the neural dynamics of baseline elevation and increased gain across different datasets. Increasing the network recurrent excitation, a parameter related to working memory, steepens the slope of accuracy log-odds as a function of RT, consistent with the empirical findings on working memory. Reducing excitatory connectivity as an analog of NMDA antagonism predicts slower and more accurate performance, consistent with the empirical causal manipulation results. Conclusion: We find that local disinhibition in a biologically plausible decision circuit with recurrent excitation and global inhibition captures a broad range of behavioral and neural SAT dynamics. These findings show that different SAT characteristics such as baseline elevation and increased gain can arise from a single circuit mechanism, highlighting the critical role of disinhibition in biological decision-making.

1-I-25 Wait wait, Don't tell me! When curiosity prioritizes the information gathering process over the outcome

Abigail Hsiung[1], Kelly Eom[1], Jia-Hou Poh[1], Scott Huettel[1], R. Alison Adcock[1]
[1]Duke University

Objective: Curiosity is a fundamental driver of information seeking. Theoretical models posit that curiosity, like reward motivation, guides choice by assigning value to information. In these models, value is derived in terms of expected reductions in uncertainty. Past studies have demonstrated that curiosity peaks approaching the most uncertain outcomes and shows a preference for early, compared to late, uncertainty resolution. However, under more dynamic environments, such as in narratives or sports, there is often a preference for maximizing uncertainty or suspense over time, suggesting that the process in which information is revealed could also confer value. In the current study, we aim to investigate, under conditions of ongoing information gathering, does curiosity favor immediate or suspended resolution? Methods: We developed a novel paradigm where participants (N = 1992, Prolific) watched videos of drawings that revealed an object over time. During the video, participants were periodically asked to report on their curiosity in the video's outcome. At the end of each report, participants were given the choice to continue watching the video or end the trial early without consequence. To isolate how curiosity guides uncertainty reduction, we manipulated the alternative choice to continuing to watch, as a between-participants condition. In our Abort-Now condition, ending

the trial early resulted in receiving no further information about the video. In our Fast-Forward condition, ending the trial early resulted in jumping to the end of the video, immediately resolving the outcome of the drawing. Results: We conducted a mixed effects logistic regression to examine how curiosity predicts the decision to stop or continue and whether this differs between our experimental conditions. Our results indicate that higher reports of curiosity significantly predicted continuation of video watching in both conditions. Surprisingly, there was a significant interaction between experimental condition and curiosity, such that under the Fast-Forward condition, curiosity was more predictive of continuing to engage with the videos than in the Abort-Now condition. Conclusion: Here, we showed that higher curiosity predicted sustaining ambiguity, rather than resolving it sooner. Our study demonstrates that curiosity not only confers value to information itself, but also influences preferences for how one arrives at that information. This finding suggests limits on the idea that curiosity drives uncertainty reduction, and points to reward value inherent in the process of discovery as an important factor in understanding non-instrumental information gathering.

1-I-26 Price heuristic hypothesis: Reconceptualizing price after a partly neuroeconomics-based theoretical study of price theory and the efficient markets hypothesis

John Haracz[1]
[1]Indiana University

Objective: According to price theory, supply and demand determine prices in markets that are at, or efficiently approaching, a self-stabilizing equilibrium (Jaffe et al., 2019). According to the efficient markets hypothesis (EMH), prices incorporate all available information and thereby signal fundamental values (Fama, 2014). This notion of "price-equals-value markets" [Statman, 2017, p. 289]) is challenged by behavioral finance and behavioral economics as well as bubbles and crashes, which suggest that prices deviate from intrinsic values. Therefore, the present theoretical study aims to reconceptualize price in a way that is more descriptively accurate. Methods: A systematic literature review focused on flaws of price theory and the EMH by using related search terms (e.g., demand curve, elasticity of demand, excess demand, disequilibrium, and limits of arbitrage). Results: Over 30 years before Tversky and Kahneman (1971, 1974) shared their insights into heuristics and biases, Hall and Hitch's (1939, p. 18) insightful survey revealed that price setters used a "full cost" "rule of thumb" rather than equating "marginal revenue and marginal cost" as in classical economic theory. This and many other examples of price-based heuristics (i.e., price heuristics) accrued before and after (e.g., price setters' "herding" [Kienzler, 2018, p. 86]) neoclassical economics came into prominence with an emphasis on price as a reliable indicator of fundamental value. However, heuristics are baked into prices (Hall and Hitch, 1939; Kienzler, 2018), which may thereby reflect price setters' biases more than intrinsic value. Accordingly, the price heuristic hypothesis proposes that prices act as heuristic signals because: a) heuristics are embedded in prices (Hall and Hitch, 1939; Kienzler, 2018); b) prices guide market participants' decision making by serving as a signal for the use of heuristics. In the bubble-crash context, traders avoiding effortful deliberation (e.g., as in neural autopilot theory [Camerer et al., 2021]) could simply use asset-price changes as guides, or heuristics, to popular or unpopular assets, thereby letting investors fall under the spell of a bubble or crash bias, respectively. These biases may promote trend-following investment strategies that yield price bubbles and crashes. Conclusions: Price theory and traditional asset-pricing models inspired by the EMH are useful tools among many other tools used by analysts and policymakers, but they are not very descriptively accurate under all conditions. The price heuristic

hypothesis is more descriptively accurate, but it has yet to be formalized and implemented as a tool for analysis and policymaking.

1-I-30 Using process models to disentangle stimulus-driven and controlled processes during value-based decision making

Harrison Ritz[1], Romy Frömer[2], Amitai Shenhav[2]
[1]Brown Unversity, [2]Brown University

Objective There is growing interest in how people use cognitive control to flexibly adapt processes involved in value-based decision-making. Grounded in motivated control research, a recent theory by Vassena, Deraeve, & Alexander (2020) proposed that people recruit cognitive control to invigorate responses when difficult decisions won't make a deadline. To support this proposal, the authors fit a drift diffusion model (DDM) to choice behavior from a speeded value-based decision making task, and show that response thresholds were lower for difficult decisions, consistent with their controlled invigoration account. Here, we critically examine the theoretical and empirical support for this proposal. Methods We re-analyzed data from this study, for which participants (n=22) chose between pairs of value-based bundles under a short response deadline (750ms). Previously, DDM modeling of these data allowed drift rate and response threshold to vary with the difference between option values on a given trial (i.e., choice difficulty). We compared the (complexity-penalized) fits of this model to alternate parameterizations of the DDM that allowed drift and threshold to vary with both overall and relative value; enabled prioritized weighting of the highest-valued items in a bundle; and/or allowed the decision threshold to collapse over time. Finally, we simulated behavior from a leaky competing accumulator (LCA) with a collapsing threshold to determine whether stimulus-driven decision dynamics in this model could mimic putative vigor-based control mechanisms. Results We found that we could substantially improve DDM model fits, relative to its previous implementation, by incorporating a collapsing bound and accounting for influences of overall set value on drift rate. When we did this, we no longer found that thresholds decreased with decision difficulty. We further show that all of the key characteristics of choice behavior on this task can be parsimoniously accounted for with a standard LCA, and that incorporating an additional vigor-based controller is both unnecessary and potentially redundant with its existing components. Conclusions Together, our results suggest that current evidence for value-based invigoration (lower thresholds for more difficult choices) is an artifact of model misspecification, and that the proposed mechanism may in most instances be obviated by existing models of value-based choice that omit such a control mechanism. More generally, our work demonstrates potential sources of mimicry between parameters of decision models that may lead to misattributions of control to evidence accumulation processes that are instead primarily stimulus

1-I-35 Theoretical models of context-appropriate adaptive learning

Leah Bakst[1], Lena Schaefer[1], Matthew Nassar[2], Joseph McGuire[1]
[1]Boston University, [2]Brown University

Objective: When a surprising outcome occurs, decision makers must determine whether it reflects meaningful change or an uninformative outlier. In dynamic environments with occasional change points, surprising events should drive substantial belief updating. Conversely, in stable environments with occasional random outliers, surprising events should be largely ignored. Previous findings show that people spontaneously calibrate patterns of adaptive learning according to the structure of their current

environment. The data also suggest an asymmetry in which people calibrate more readily to change-point environments than to environments that require ignoring outliers. Here, we developed and tested theoretical models to account for these phenomena. Method: We developed two models of behavior in a gaze-based implicit spatial prediction task. The task had two conditions: one in which large event-to-event position changes tended to reflect change points (CP) in the generative process, and one in which the generative process followed a random walk (RW) interrupted by occasional uninformative outliers. The first candidate model tracked the predictions of two approximately Bayesian models specifically tuned to the CP and RW conditions and accumulated event-by-event evidence to infer which condition was in effect ("two-alternative model"). A second model lacked foreknowledge of either environment ("agnostic model") and generated predictions by flexibly inferring the informativeness of outcomes as a function of their associated prediction error magnitude. Results: The two-alternative and agnostic models predicted actual target locations less successfully than a model with built-in knowledge of the true environmental structure but still significantly better than participants (all $p < .001$, Wilcoxon signed-rank test). The two models showed comparable levels of success in predicting behavior ($p = .581$), with the agnostic model better mimicking participants' tendency to become better calibrated over time. Internal dynamics of the two-alternative model showed that the strength of evidence for the two conditions was asymmetric: cumulative evidence strength was significantly greater for CP than for RW ($p < .001$). Conclusion: We developed two models that calibrated adaptive learning without advance knowledge of environmental structure. The models provide a window into how people adopt context-appropriate patterns of adaptive learning from feedback. The observed asymmetry between conditions offers a potential explanation for previous findings of better-calibrated behavior in environments where extreme events are meaningful than where they should be disregarded.

1-I-9 Information-seeking during COVID-19

Nathan Torunsky[1], Kara Kedrick[1], Iris Vilares[1]
[1]University of Minnesota - Twin Cities

Objectives: In the context of a global pandemic, understanding the motivations behind information-seeking can provide useful insights for public health education and intervention. Sharot & Sunstein (2020) recently proposed that information-seeking is driven by estimates of the hedonic, instrumental, and cognitive utility of potential resources. Additionally, past research has shown that information-seeking can be driven by emotional states and tendencies. The goal of the present study was to identify predictors of COVID-19 information-seeking and examine the relationship between individual differences in emotion and the calculation of information value. Methods: We recruited 187 undergraduate participants from the University of Minnesota during the fall of 2020. Participants completed five rounds of an information-seeking task in which they read about a topic related to COVID-19, rated their expected hedonic, instrumental, and cognitive utility of an article excerpt on that topic, and decided whether to seek or avoid the excerpt. After finishing the information-seeking task, participants completed self-report questionnaires assessing their general mood state (i.e. depression, anxiety, and stress), intolerance of uncertainty (IU), engagement in COVID-19-preventative behavior (e.g. mask-wearing), and trust in science and scientists. In addition to Spearman correlations, we fit mixed-effects logistic regression models to predict participants' information seeking behavior, and used BIC and cross-validation to select the best-fitting model. Results: Model selection yielded a model with only a single predictor of information-seeking: cognitive utility ($\beta = 1.2$ [0.90, 1.6], $p = 2.2e\text{-}11$, OR = 3.5 [2.5, 5.1]). Depression, anxiety, and IU were not correlated with overall information-seeking or

participants' estimates of information utility, but preventative behaviors positively correlated with information-seeking (r = .21, p = .006). Trust in science and scientists was positively correlated with information-seeking, but failed to reach significance, possibly due to poor variation in scores (r = .14, p = .064). Trust in science and scientists was strongly correlated with engagement in COVID-19 preventative behaviors (r = .31, p = 3.3e-5). Conclusions: In an undergraduate sample, information-seeking about COVID-19 appears to be primarily driven by the expected cognitive utility yield of a particular resource rather than by hedonic or instrumental utility or individual differences in emotional states. Additionally, self-reported preventative behaviors appear to correlate both with information-seeking and trust in science and scientists.

1-J-27 From valence to value: Neurocomputational mechanisms for transforming affect into decision-relevant information

Ian Roberts[1], Azadeh HajiHosseini[1], Cendri Hutcherson[1]
[1]University of Toronto

How do emotions determine the value of a stimulus during choice? Intuitively, pleasant affect should make a stimulus more valuable, which might suggest the two terms are interchangeable. Yet a growing body of research demonstrates that people often value objects precisely because of the unpleasant feelings they elicit. These insights raise an important question: if affect and value are not synonymous, when and how does the brain transform one into the other? Are extra processes required to transform unpleasant affect into positive value? What differentiates people who more flexibly value different emotions? To address these questions, we developed a neurocomputational model of affect valuation in which people convert subjective affect into context-sensitive decision value through a process of weighted evidence accumulation. We then tested model predictions by recording EEG and facial EMG during a novel affective choice paradigm in which affective valence and decision value are orthogonalized. Our model fit shows that people readily adapt to different affect-valuation contexts through a simple re-weighting of affective inputs to evidence accumulation processes - finding little evidence for additional modifications between contexts. Furthermore, supporting model predictions, single-trial regression analyses revealed early affect-related neural signals emerging in parietal (~250ms) and frontal (~400ms) channels followed by a later value-related signal in frontocentral channels (~500ms). The temporal transition from emotion generation to emotion valuation processes closely aligned with our model's estimated onset of evidence accumulation (507ms). Moreover, we find that individual variability in our computational estimates of flexible affect valuation across contexts is associated with modulations to neural signals associated with emotion valuation but not emotion generation. This suggests that people who exhibit reduced affect valuation flexibility do so because of more rigid affect-value associations rather than either amplified or muted emotions. Finally, while EEG suggested a temporal progression from affect to value, facial EMG signals from corrugator and zygomatic muscles showed sustained representations of affect (in both) and value (in zygomatic) throughout the choice process suggesting that not all affective indices show the same transition to value. In sum, our findings introduce and support a new computational model of affective choice and support the notion that affect and value are different constructs.

1-J-28 Culture selectively shapes neural responses to social rewards

Elizabeth Blevins[1], Michael Ko[1], BoKyung Park[2], Yang Qu[3], Brian Knutson[1], Jeanne Tsai[1]
[1]Stanford University, [2]University of Texas at Dallas, [3]Northwestern University

Prior research has demonstrated that European Americans ideally want to feel excitement and other high arousal positive affective states more than Chinese. Consistent with cultural differences in "ideal affect," European Americans showed greater neural responses in regions associated with reward (e.g., the Nucleus Accumbens or NAcc) when viewing excited versus calm smiling faces compared to Chinese. It remains unclear, however, whether these cultural differences are specific to social incentive stimuli rather than more general incentive stimuli, and whether these differences have implications for peoples' real-world social behavior. To probe these questions, healthy European American (N = 26) and Chinese (N = 27) adults played the Social Incentive Delay (SID) task, in which they were cued to anticipate and receive social rewards (i.e., smiling faces) that varied in magnitude (calm, moderate, excited). Participants then played the Monetary Incentive Delay (MID) task, in which they were cued to anticipate and receive monetary rewards that varied in magnitude ($1, $3, $5). After scanning, participants rated the emotional expressions of six friends in their social media profile photos on a scale that ranged from -3 (most intense negative expression) to 3 (most intense smiling expression). As predicted, European Americans showed greater NAcc activation than Chinese while viewing excited smiling faces in the SID task, B = .09, SE = .04, t(703) = 2.50, p = .01 (European American: M = .19, SE = .02; Chinese: M = .10, SE = .02). In contrast, there were no significant differences between European Americans and Chinese in response to gaining $5 in the MID task, p = .75 (European American: M = .07, SE = .03; Chinese: M = .09, SE = .02). Moreover, across cultures, the greater participants' NAcc responses to excited smiling faces, the more intense their friends' smiles were in their social media profile photos (r = .30, p = .03). These findings suggest that culture can selectively shape neural responses to social rewards, and that these responses are associated with individuals' social relationships in everyday life.

1-K-29 Gaze dynamics in many-option choice

Wenjia Joyce Zhao[1], Armin Thomas[2], Sudeep Bhatia[3], Ian Krajbich[1]
[1]The Ohio State University, [2]Stanford University, [3]University of Pennsylvania

Attention is a key determinant of value-based choice. Yet we currently lack a general quantitative framework capable of providing a systematic account of attentional dynamics in large and complex choice sets, such as those encountered by decision makers in everyday choice situations (e.g. when choosing products in a grocery store). Here, we build such a framework and apply it to eye-tracking data from a many-option food choice experiment. Our approach is based on established theories of attention and memory, and describes nuanced aspects of visual search dynamics, i.e. where people look at a given point in time and how this depends on where people have looked previously. In the experiment, each subject (N=49) completed 200 incentivized food choice trials. In each trial, K options (K in {9, 16, 25, 36}) were displayed in a square grid, and choices and eye-movements were recorded. Following the choice task, subjects rated the items based on their willingness to eat them. To predict the probability of looking at (i.e. "sampling") each item in each gaze, we distinguish among three types of factors. Type-1 factors are independent of all previous samples (e.g., distance to the center of the screen). Type-2 factors depend only on the most recent sample, and account for things like spatial proximity and reading-order. Type-3 factors involve a (potentially decaying) influence from all previous samples, with decreased sampling of (recently) sampled items and increased sampling of items with high values. We instantiate these factors as parameters within a dynamic and stochastic modeling framework designed to predict gaze patterns. We fit numerous model variants for each subject and set size, using Bayesian techniques. We find strong evidence for all the factors in the observed gaze patterns, but some heterogeneity across subjects. Importantly, our approach quantitatively predicts key properties of the

gaze patterns in the data such as the high probability of sampling neighbors, the frequent sampling and resampling of high-value items, and the delays before returning to an item. Overall, our quantitative, tractable, and general modeling framework provides novel insights regarding visual search dynamics in complex, many-option, value-based choice. In doing so, it allows for the study of difficult but intriguing research questions regarding the interactions between attention, value, and choice in everyday decision settings.

Poster Session #2

2-B-41 Neural responses clarify how eco-labels can promote energy-efficient purchases

Tara Srirangarajan[1], Nik Sawe[1], Anshuman Sahoo[1], Grace Tang[1], Brian Knutson[1]
[1]Stanford University

Although consumers have historically purchased inefficient appliances with high long-term energy costs, behavioral and policy interventions such as ecolabels (e.g., the Energy Star label) have offset this trend by increasing energy-efficient purchases. While the Energy Star label was designed to prompt consumers to evaluate trade-offs between immediate prices and future energy savings, it may more simply increase products' perceived attractiveness. To reveal the mechanisms underlying the effectiveness of the Energy Star label, we combined incentive-compatible behavioral experiments, neuroeconomic measures, and a national stated choice survey. In a Functional Magnetic Resonance Imaging (FMRI) experiment (n=36), the Energy Star label increased activity in neural regions associated with positive affective responses that promote purchasing (i.e., the Nucleus Accumbens or NAcc), particularly in individuals with steep delay discounting rates (e.g., more financially impulsive). Further, group-averaged neural data (from the NAcc) was associated with consumer demand in a large national stated choice experiment (n=1,550). These results suggest that ecolabels leverage affective responses implicated in purchasing at both individual and aggregate levels rather than facilitating cognitive assessment of trade-offs between initial price and long-term efficiency. Ultimately, these findings may help policy makers refine ecolabels to further increase investment in energy efficiency and optimize behavioral interventions.

2-D-12 Late positive potential as a neurophysiological marker of evidence accumulation in social and non-social categorization

Hongbo Yu[1], Sai Sun[2], Jinge Wang[3], Xin Li[3], Shuo Wang[4]
[1]University of California Santa Barbara, [2]Tohoku University, [3]West Virginia University, [4]Washington University in St. Louis

Objective: Recent studies have shown that a late positive potential (LPP) in the EEG signal tracks decision speed in both social and non-social categorization tasks. However, the computational role of the LPP in resolving ambiguity in categorization is unclear. Here, by applying drift diffusion models to three EEG datasets (total N = 108), we showed that the LPP reflects an evidence accumulation process in resolving both social and non-social categorical ambiguity. Methods: Participants viewed static images while undergoing EEG recordings. In Study 1, three groups of participants were randomly assigned to judge fear-happy morphed faces (N = 23), anger-disgust morphed faces (N = 11), or cat-dog morphed images (N = 11). In Study 2, three groups of participants viewed fear-happy morphed faces, but with different tasks - fear/happy judgment (N = 16; morph level was task relevant), female/male judgment (N = 16; morph level was task irrelevant), or poor/rich judgment (N = 16; morph level was task irrelevant).

In Study 3 (N = 15), participants judged whether anger-disgust morphed faces and cat-dog morphed images were faces or animals. Thus, the morph level (i.e., perceptual ambiguity) was task irrelevant. We estimated drift diffusion models (DDMs) and examined whether the drift rate was modulated by the interaction between stimulus category (e.g., fearful vs. happy) and trial-by-trial LPP magnitude. We hypothesize that if LPP reflects an abstract decision signal, then the drift rate should be higher when LPP magnitude is larger, but this will only hold when ambiguity is task relevant. Results: In Study 1, although different groups viewed different types of stimuli, the perceptual ambiguity was equally relevant to their judgments. Here, the interaction between stimulus category and LPP magnitude significantly modulated drift rate, and the modulatory effect did not differ across tasks, indicating that the LPP reflects the speed of evidence accumulation in categorization. In Study 2, the modulatory effect of the interaction term on drift rate was significantly stronger in the task relevant group than the task irrelevant groups. In Study 3, when the perceptual ambiguity was task-irrelevant, the interaction term did not significantly modulate the drift rate. Conclusions: By applying DDM to model EEG signals, we demonstrated that the LPP reflects an evidence accumulation process in both social and non-social categorization. Critically, the evidence accumulation process that the LPP captured was not specific to the content of the stimuli (e.g., social or non-social), but was sensitive to task demand, thereby reflecting a domain-general decision signal.

2-D-42 ¬Controlling ambiguity: The illusion of control in decision-making under risk and ambiguity

Alex Berger[1], Agnieszka Tymula[1]
[1]University of Sydney

The finding that individuals are ambiguity averse in the domain of medium likelihood gains is robust across laboratory studies. We conducted a laboratory experiment to examine whether we can lower this aversion to ambiguity by giving participants an illusion of control over the probability of winning by allowing them to choose the winning color in what objectively are 50-50 gambles. We found that while our illusion of control manipulation does not increase participants' preference for risky gambles with a known probability of winning, it does increase participants' tolerance for ambiguous gambles with unknown probability of winning. When the illusion of control is absent, the structural model estimates of ambiguity tolerance are 29% lower. Our results highlight the importance of considering the illusion of control in the estimation of ambiguity attitudes.

2-D-43 The Impact of Background Ambiguity On Risk Taking

Agnieszka Tymula[1], Vinayak Alladi[1], Mahdi Akbari[1]
[1]University of Sydney

Although it has been established, empirically and theoretically, that background risk impacts decision-making, no studies have examined the impacts of background ambiguity (rather than risk) on risk-taking. We present results from a within-subject experiment in which subjects made a series of investment decisions by allocating a portion of their controlled wealth towards a risky asset in the presence of various forms of background uncertainty (both risk and ambiguity). Our main result is that 44% of our subjects showcased some form of background ambiguity aversion, defined as being more averse to background ambiguity than to background risk. In the aggregate, subjects invest approximately 6 p.p. less on average in the risky asset when there was background ambiguity as opposed to background risk. In addition, we find that contrary to Gollier and Pratt's (1996) theoretical work on background risk,

individuals classified as having CARA (constant absolute risk aversion) and DARA (decreasing absolute risk aversion) are also background risk-averse. Finally, we use variation in income and socioeconomic status and find evidence that poor individuals react less to background uncertainty (both risk and ambiguity) than the rich.

2-D-44 Frontal but not parietal cortex is required for economic decisions under risk

Xiaoyue Zhu[1], Joshua Möller-Mara[1], Sylvain Dubroqua[1], Chaofei Bao[1], Jeffrey Erlich[1]
[1]NYU Shanghai

Objective: Neural activity in the posterior parietal (PPC) and frontal premotor cortex has been shown to correlate with economic decision-making. However, to our knowledge, the causal role of these areas in economic choice has not been tested. Methods: Here, rats earned rewards in a novel risky-decision task, where, on each trial, they were offered both a small certain reward (sure-bet) and a lottery whose size was indicated by an auditory cue. Importantly, the lotteries were randomized and independent, thus the animals made cue-driven choices on a trial-by-trial basis. Eight rats were implanted with bilateral cannula in the PPC and in the frontal orienting field (FOF). Results: PPC silencing with muscimol caused minimal effects (−1.17±0.63, p=0.06, GLMM), but FOF inactivation significantly decreased lottery choices (−2.68±1.03, p <0.01). We decomposed the effect of FOF inactivation on behavior using a multi-agent model that included a 'rational' utility-maximizing agent and two 'habitual' agents. While silencing of the FOF resulted in a decrease in lottery choices, the risk-preference of the rational agent was surprisingly unchanged. Instead, the influence of the rational agent was diminished (95% C.I. = [0.37 0.49] vs. control 0.87,p <0.001), and a 'habitual' sure-bet agent dominated the final decision (95% C.I. [0.51 0.62] vs. control 0.06, p < 0.001). Consistent with previous studies, unilateral PPC silencing produced a marked ipsilateral bias (T5= 3.09, p <0.05) on 'side-choice' trials where the rewards were equal, providing positive control for the null results in risky trials. Finally, PPC inactivation did not impair the ability to learn new sure-bet values in 4 rats. Conclusions: It was found that FOF perturbation shifted choices from a goal-directed mechanism to a habitual one, without changing the parameters of the rational agent. This suggests that FOF is not the substrate for the rational agent, but it is a key bottleneck by which the rational agent can influence the subjects' final choice. Our results also show that PPC plays a minimal role in not only economic choice under risk, but also the learning of new reward values. To our knowledge, our study provides the first causal evidence of the role of the frontal-parietal circuit in economic choice in rats, or any other species.

2-D-45 Discovering the temporal correlation between environmental states and risk preferences

Zih-Yun Yan[1], Paul Glimcher[1]
[1]New York University

Environmental states affect our day-to-day decisions, including our choices under risk and uncertainty. Both human and animal experiments showed that the internal metabolic states could alter the risk preferences. (Symmonds et al., 2010; Levy, Thavikulwat & Glimcher 2013; Yamada, Tymula, Louie & Glimcher 2013). To extend these findings from the internal state to the external activity states at a broader time scale, we performed longitudinal examinations of daily activities and risk preferences to gathered the within-subject data. We aimed to identify how the engagement of different activities could predict the temporal change of the risk preferences. All subjects were recruited in the US on CraigsList and NYU subject recruitment system. 115 out of 123 subjects in total finished the experiments. We

employed a mobile smartphone-based experimental platform (Linkt) for gathering daily data from our participants. Each participant was asked to report time use diary per week over a two-month window and the following instruments at least once every week: a risky choice task (Levy et al. 2010), a delay discounting task (Kable & Glimcher, 2007), the self-report positive and negative affect (Kahneman et al. 2004). In our activity tracker, we asked subjects to fill out a timeline of everything they did on that day, and we categorized these activities into sleep, work, home, leisure, and social times. In addition to the temporal measurements, we also gathered 13 personality inventories throughout the study from the following domains: attitude toward uncertainty, temporal discounting tendency, impulsivity, and the level of psychological distresses. Here we present data showing how the amount of a particular activity spent last week correlates with the risk attitude-related variable. For the risk measurement, we used PCA analysis to find the principal components that risk-related measurements from the non-parametric and parametric method can explain the variance the most and use it as the single risk score. We next use these risk scores and the activity time to perform a time-series autoregression model. We found that the sleep time from last week can significantly predict the risk score in the current week, suggesting that with more sleep time, subjects become more risk-tolerant. We did not find a significant predictive effect from the other types of activities. For the next step, we will expand our data collection to a longer time window to capture more dynamics from each subject. Disclosure: PWG is an officer and stockholder in Datacubed Health.

2-D-46 Decisions under risk in mice, rats, and humans

Joshua Moller-Mara[1], Xiaoyue Zhu[1], Sylvain Dubroqua[1], Evgeniya Lukinova[1], Jeffrey Erlich[1]
[1]NYU Shanghai

Objective: All species make decisions under risk. Here, we examined the strategies employed by a large cohort of mice, rats and humans in a non-verbal risky-choice task that emphasized known aleatory risk: that is, risk was not due to Knightian uncertainty. Our goal was to determine the cross-species similarities and differences in decision making under risk. Having a single well-validated behavioral task that could be employed across species would lay a foundation for circuit level investigation of neural mechanisms of risky-decisions. Methods: We trained mice, rats (n = 39) and humans (n = 37) in a novel non-verbal risky-decision task. On every trial, an auditory cue communicated the value of an all-or-nothing lottery (the probability of winning remained fixed). Subjects chose between that trial-specific lottery or a fixed small-but-certain reward. The audio cue was randomized across trials, so a utility-maximizing subject would use the cue to guide their decision and avoid any trial-history-based strategies. To assess the strategies used by subjects of each species, we fit choices using Bayesian mixture models that incorporated classic economic elements (like marginal utility and variance aversion) with behavioral elements more commonly used in models of animal behavior, like choice-bias and history-dependent strategies (e.g. win-stay / lose-shift / perseverance). Toggling these elements created a family of 15 models which we fit to each subject and used 10-fold cross-validation for model selection. Results: Our smallest model, a classic economic 2-parameter marginal utility model (utility curvature, ρ, and noise, σ) provided the worst fit to the data for all three species. The best models all included a "rational" utility-maximizing agent with ρ, σ, and a variance penalty, κ, mixed with "habitual" agents, which chose an action ignoring the current-trial cue (the mixing fraction, ω, determined the relative influence of the rational and habitual agents on choice). The main difference between species at the model-selection-level was that the fit of human subjects' choices was not improved by the addition of trial-history effects, whereas fit of the rodent subjects was improved by addition of history-effects.

Using the wrong model caused a bias in estimation of ρ. For the compact ρ-σ model, we estimated ρ to be 0.98 ± 0.22, 0.63 ± 0.11, 0.65 ± 0.15 (μ ± s.d.) for humans, rats and mice respectively. When we used the best model we estimated ρ to be 1.05 ± 0.26, 0.86 ± 0.13, 1.09 ± 0.19. Conclusions: Decomposing choices under risk requires a mixture model. The model shows differences (novel insights) in species on the same task.

2-E-36 Deliberative evaluation in intertemporal choice is shaped by experiment structure

Nidhi Banavar[1], Aaron Bornstein[1]
[1]University of California - Irvine

Objective: Modeling choice behavior or response time in randomized, or unstructured, intertemporal choice (ITC) experiments often assumes an absence of trial-level sequential dependencies. A natural extension of popular theories of ITC that involve simulating the future (Peters & Buchel, 2010) imply that already computed future values could be cached and re-used, especially if an individual makes similar choices in sequence (Dasgupta et al, 2018). Here we test how sensitive our inferences about the components of deliberative evaluation are to such "spillover" effects through response time modeling. Methods: Subjects (n = 482) participated in a text-based ITC task with no explicit sequential dependencies (Hunter, Bornstein, Hartley, 2018). Each subject made a sequence of 102 incentive-compatible decisions where they chose between a smaller amount of money today or a larger amount of money in the future. To link response time modeling with stimulus properties, we developed a Bayesian implementation of a Drift Diffusion Model approximation (Bogacz et al, 2006). We then applied a hierarchical framework that allows for distinguishing whether response times are influenced by spillover between ostensibly unrelated trials. This framework can simultaneously discern whether there is meaningful evidence for spillover on multiple different aspects of the cognitive processes under examination - e.g. on discount factor, drift rate, or bias. Results: Nearly one in three subjects show meaningful evidence (BF > 3) of sequential effects in the drift rate and/or overall bias towards patience or impulsivity. These effects are driven by both cross-trial value and delay differences, as well as their interaction. Approximately half of these effects manifest on an individual's bias parameter, to the degree that parameter interpretation changes: sequential effect adjusted parameters show individuals inferred to be patient (prefer Later option) in the standard approach as actually having an overall preference for the immediate option. Our model also successfully captures a wide range of individual differences in both the presence and absence of spillover in a non-sequential task. Conclusions: Our results highlight the importance of explicitly accounting for the fact that humans process information sequentially. We demonstrate that when behavioral data is analyzed with this fact in mind, the estimated variables in higher order cognition can change not only in magnitude but also in interpretation. This is critical for improving the mapping between parametric estimates and real-world behavior.

2-E-40 Investigating the link between neural reward reactivity and attention

Nitisha Desai[1], Allison Londerée[1], Eunbin Kim[1], Dylan Wagner[1], Ian Krajbich[1], Kentaro Fujita[1]
[1]The Ohio State University

Self-control failures--the undermining of long-term goals for short-term rewards--are implicated in many societal problems, including obesity, substance abuse, and poor financial decisions. Previous research has shown that self-control failure is related to greater activity in the brain's reward regions (e.g.,

ventral striatum) when exposed to temptations. By contrast, greater activity in the regions involved in inhibitory control (e.g., inferior frontal gyrus) supports successful resistance to temptation. Along with these neural mechanisms, visual attention has also been shown to play a role in self-control decisions. Research has found that attentional biases towards short-term rewards are associated with self-control failures. Here--using a multi-modal approach that combines functional MRI (fMRI), eye-tracking, and diffusion modeling--we investigate the link between neural reward reactivity to appetitive stimuli and attentional biases to short-term rewards. In our experiment, 26 dieters participated in an fMRI scan followed by an eye-tracked choice task. Dieting status was confirmed by commonly used metrics of dieting concern (i.e., Revised Restraint Scale). We used the scanning session to measure individual differences in reward cue-related activity to appetitive foods. After the fMRI session, we assessed subjects' self-control by having them make binary choices between healthy and unhealthy foods. There were 300 trials in total; 100 "conflict" trials directly pitted a healthy food against an unhealthy food. We found that activity in the left nucleus accumbens (NAcc) predicted the tendency to look at the unhealthy item first ($r(24)=0.55$, $p=0.003$), whereas activity in the inferior frontal gyrus (IFG) was related to the probability of looking at the healthy item first ($r(24)=0.39$, $p=0.047$). To further explore the relationship between reward reactivity and attention, we fit an attentional drift-diffusion model. Modeling results revealed that activity in the right ventromedial prefrontal cortex activity was negatively correlated with starting point bias (z) towards the healthy option ($r(24)=-0.41$, $p=0.038$) and with the attentional discount (θ) on the not-looked-at option ($r(24)=-0.46$, $p=0.017$). Our results suggest greater NAcc activity may be associated with a bias to initially seek out short-term rewards, while greater IFG activity may be associated with a bias to seek out the options consistent with long-term goals. These novel findings provide a first step towards identifying a mechanism--early attentional biases--by which neural reward reactivity impairs self-control and neural inhibitory activity promotes it.

2-E-47 Intertemporal choice reflects value comparison rather than self-control: insights from metacognitive confidence

Adam Bulley[1], Karolina Lempert[2], Colin Conwell[1], Muireann Irish[3], Daniel Schacter[1]
[1]Harvard University, [2]University of Pennsylvania, [3]University of Sydney

Objective: Intertemporal decision-making between sooner and later rewards has long been assumed to measure self-control, with prominent theories treating choices of smaller, sooner rewards as a failure of attempts to override immediate temptation. If this view is correct, people should be more confident in their decisions when they "successfully" delay gratification. Here, we tested this assumption and charted how decision-makers judge which intertemporal choices are the right ones for them. Methods: In two pre-registered online experiments with built-in replication, adult participants (n=117) made a series of monetary intertemporal choices and rated their confidence in having made the right decision after each choice. Results: Contrary to assumptions of the popular self-control account, confidence was not reliably higher (or lower) when participants chose delayed rewards ($B = 0.01$, $p = 0.904$). Rather, participants were more confident in their decisions when available rewards were further apart in subjective value ($B = 0.13$, $p < 0.001$), less delayed ($B = -0.08$, $p = 0.005$), and larger in magnitude ($B = 0.04$, $p = 0.011$). Demonstrating metacognitive insight, participants were also more confident in decisions that lined up closer with their independent valuation of possible options ($B = 1.32$, $p < 0.001$). Decisions made with less confidence were more prone to changes-of-mind ($B = -0.76$, $p < 0.001$). and more susceptible to a patience-enhancing manipulation ($B = -0.42$, $p < 0.001$). Conclusions: Together, these results establish that metacognitive confidence in intertemporal choice tracks uncertainty in value comparison - just as it

does in other kinds of decisions that are unrelated to self-control. Our findings thereby challenge self-control views and instead support the idea that intertemporal choice reflects value comparison between sooner and later reward options.

2-E-48 Quasi-hyperbolic present bias: a meta-analysis

Xueting Wang[1], Stephen Cheung[1], Agnieszka Tymula[1]
[1]Univeristy of Sydney

Quasi-hyperbolic discounting is one of the most well-known and used models to capture self-control problems. The underlying assumption of the model is that agents have a "present bias" toward current consumption, as all future rewards are downweighed relative to rewards in the present, in addition to the standard exponential discounting of delayed rewards. To create a meta-analytic dataset of present bias estimates, we searched all major research databases using the intersection of a set of methodology keywords with a set of topic keywords. We identified 62 papers with 81 estimates of present bias parameter. The literature shows that people are present biased towards money (β=0.82), but substantial heterogeneity across studies exists. The source of heterogeneity comes from the subject pool, methodology (e.g. BDM auction), geographical location of data collection, payment method, the study place (e.g. online vs. lab). Reward type also has an influence on the estimates of present-bias parameter: individuals show stronger present bias towards real effort and health outcomes compared to monetary rewards. There is evidence of selective reporting and publication bias in the direction of overestimating present-bias, but the present-bias still exists after correcting for these issues.

2-E-49 Reaction Time and Impulsive Decision Making in a Delay Discounting Task

Lawrence Amsel[1], Diana Rodriguez-Moreno[1], Xiaoxiao Sun,[1], Christina Hoven[1], Yael Cycowicz[1]
[1]Columbia University

It has been suggested that impulsivity can be measured in delay discounting (DD) tasks as the preference for smaller sooner (SS) reward over larger later (LL) rewards. We investigated the relationship between the decision of choosing SS over LL rewards and reaction time (RT). In perceptual decision making, as the difference between the stimuli gets larger the choice becomes cognitively less demanding and the RT shorter. In contrast, we tested the hypothesis that an impulsive cognitive approach generates SS choices, which are thus made faster - shorter RT- regardless of the reward differences. METHODS: Participants (N=20; age: 22.5±4.0; female=7) performed a DD task (on average 124 trials). in an fMRI scanner (MRI not discussed here). We included a reward-immediacy condition (SS reward was available either at the time of task or in 2 weeks) and a framing condition (the default choice was either SS or LL). The data were analyzed at the trial level without aggregation. A Bayesian generalized linear mixed-effects model was tested on SS vs. LL choices as a function of reward immediacy, framing, SS and LL time difference (2 or 4 weeks), SS reward magnitude, SS and LL relative reward difference (RRD), and RT. We also included an interaction terms for relevant factors with RT, and a random intercept varying over participants. RESULTS: A significant main effect of RT was found, indicating that with increased RT participants made more LL choices. In addition, RT interacted with RRD, demonstrating that when the choices had smaller reward differences shorter RT was associated with fewer LL responses, whereas when the choices had the largest RRD the probability of LL responses were not associated with RTs. We performed the same analysis with Group as an additional factor based on individuals' prior performance on the same task, with those who made more SS choices were the

reward impulsive group (N=10), and the reward patient group included individuals who made fewer SS choices. We found that RT interacts with Group. The reward impulsive group showed a steeper increase of LL choices with increased RT compared to patient participants. CONCLUSION: These data demonstrate that, unlike perceptual decision making, RT is related not only to task difficulties but also to impulsivity. While longer RT was associated with a gradual increase in LL choices with increased RRD, shorter RT demonstrated a steeper increase of LL choices with increasing RRD. Also, impulsive, compared to patient individuals, showed a steeper change to LL choices with increasing RT, suggesting that the contribution of cognitive processes involved in LL choices depends to some degree on impulsivity.

2-F-50 Connectome-based Individualized Prediction of Reciprocity Behavior in Give and Take Frames

Huihua Fang[1], Chong Liao[1], Fu Zhao[2], Shuang Tian[2], Yuejia Luo[2], Pengfei Xu[3], Frank Krueger[4]
[1]Department of Psychology, University of Mannheim, Mannheim, Germany, [2]Shenzhen Key Laboratory of Affective and Social Neuroscience, Magnetic Resonance Imaging Center, Cen, [3]Beijing Key Laboratory of Applied Experimental Psychology, Faculty of Psychology,

Reciprocity --a social preference intrinsically possessed in human nature that one tends to return the act given by others-- forms the pillar of human interpersonal relationships. Several neuroimaging studies have recently investigated the neuropsychological signatures behind the heterogeneity of reciprocity. For example, whole-brain resting-state functional connectivity (RSFC) and within-network RSFC of cingulo-opercular network (CON), default mode network (DMN), and frontoparietal network (FPN) significantly predict reciprocity. Despite its stable trait property, evidence exists that reciprocity is susceptible to social framing, i.e., different descriptions of an identical outcome may greatly influence one's social decision. People, for example, reciprocate more in the "give" frame as compared to the "take" frame although the objective outcomes are equivalent in both frames. However, it remains unknown whether RSFC can predict reciprocity described as "give" or "take" frames as well as the framing effect in reciprocity. We collected RSFC fMRI data and reciprocity decisions from ninety healthy volunteers, who completed one-shot versions of the trust ("give" frame) and distrust ("take" frame) game. First, we analyzed the difference of reciprocity between "give" and "take" frames as well as the association between both frames. Second, we performed a linear support vector regression analysis, combining a leave-one-out cross-validation with a permutation approach, to predict individual differences in reciprocity for the two frames and the framing effect (the difference of reciprocity rate between the two frames) based on RSFC (Dosenbach atlas: DMN, FPN, CON, sensorimotor network [SMN], occipital network [OccN]). Our behavioral results showed that participants reciprocate significantly more in the "give" compared to the "take" frame and the reciprocity rate for both frames was positively correlated. Our neural results revealed that whole-brain RSFC predicted both reciprocity in the "give" and "take" frame and the framing effect. The prediction of reciprocity in the two frames showed a similar pattern: CON (probably associated with control processes for strategic behavior, i.e., applying different strategies based on partners' action) and DMN (likely associated with mentalizing, i.e., inferring a partners' intention and expectation) outperformed other networks in terms of the total number of predictive edges, and all networks contributed comparably in the prediction of the framing effect. In summary, our results shed light on how large-scale intrinsic brain networks are contributing to the explanation of people's social behavior of reciprocity.

2-F-51 Exogenous melatonin affects behavioural reactions to unfair offers in the Ultimatum Game

Misa Kurihara[1], Hideki Ohira[1], Pranjal Mehta[2]
[1]Nagoya University, [2]University College London

Emotion regulation is important for better social relationship, but sometimes, people fail to regulate emotions and emotionally revenge others. Recently, melatonin, a sleep hormone, has been considered to be involved in antisocial behaviours. Both animal and human studies suggest that melatonin increases antisocial behaviours such as reactive aggression (Liu et al. 2017; Jasnow et al. 2000), suggesting that melatonin may increase emotional process. In this study, we explored how melatonin contribute to emotion regulation in the social decision making using the Ultimatum Game. This study was cross-over, double-blind, and placebo-controlled experiments. Participants took part in the two sessions, sparing the days between 3 and 7 days between the two. In this study, participants were orally administered a 2mg tablet of melatonin or placebo and played the Ultimatum Game as a role of responder. After completing the Ultimatum Game, they answered a questionnaire about the Ultimatum Game. Their emotions and fairness perceptions to each offer size were reported in the questionnaire. As we predicted, exogenous melatonin increased the selection of "reject" in the Ultimatum game. While it was statistically significant when the offer was extremely unfair (£1 vs £9), we could see the same pattern to each offer size in that participants rejected the offers more in melatonin condition than placebo condition. Melatonin administration did not affect sleepiness, mood, fairness perceptions and emotions to unfair offers from self-reports, but melatonin increases the rejection when melatonin sleepiness is high and melatonin slows down reaction times in the UG. The selection of rejection in the UG has been explained by several psychological reasons such as altruistic punishment, retaliation and reward processing. Given that melatonin increases rejection when sleepiness is high, and that melatonin increases response time in the UG, melatonin may affect emotion control. In this study, we might not be able to observe the change in their emotions to unfair offers in the UG from their self-reports because detecting emotions from their self-reports is difficult (Hubbard, et al, 2002). In the future, it is expected to see how melatonin affects behavioural reactions to unfair offers and if melatonin modulates emotion and emotion regulation using additional psychophysiological measurements, tasks and fMRI.

2-F-64 Harm valuation in moral decisions shapes individual differences in two dimensions of utilitarianism

Yoonseo Zoh[1], Hongbo Yu[2], Luis Sebastian Contreras-Huerta[3], Annayah M. B. Prosser[4], Michel-Pierre Coll[5], Matthew A. J. Apps[6], Geoffrey Bird[3], Walter Sinnott-Armstrong[7], Steve W. C. Chang[1], Molly Crockett[1]
[1]Yale University, [2]University of California, Santa Barbara, [3]Oxford University, [4]University of Bath, [5]McGill University, [6]University of Birmingham, [7]Duke University

The moral principle of utilitarianism states that the morality of actions depends on their consequences. Recent work suggests that individual differences in utilitarian tendencies fall along two dimensions: a permissive attitude towards harming others (instrumental harm, IH) and maximizing welfare impartially (impartial beneficence, IB). Here, we investigated how IB and IH relate to the valuation of consequences for self and other using a task where participants traded different amounts of money for self against different numbers of painful electric shocks delivered to either themselves or an anonymous stranger. In studies of both incentivized (N=80) and hypothetical (N=192) choices, IB predicted a lower likelihood of choosing harmful choices for others relative to the self, while IH predicted a greater likelihood of choosing to harm others for one's own monetary benefit. A computational model of choice behavior,

which captured participants' valuation of pain for self and others, showed that these associations arise from individual variability in computation of the distinct values of harm to self and other. Moreover, applying a computational model of the value-accumulation process driven by weights on choice attributes revealed that IB was associated with an increased weighting of pain during decisions about harming others relative to self, whereas IH was associated with a decreased weighting of pain across decisions about harming both self and others. Neurally, we found that the similarity in two dimensions of utilitarian tendencies between participants, even after controlling for shared variance in harm aversion, predicted increased similarity of parametric responses to other's pain in prefrontal cortex and subcortical areas. Our results reveal distinctive behavioral preferences and shared neural processing associated with two dissociable aspects of utilitarianism, and the role of self- and other-regarding harm valuation in shaping utilitarian moral tendencies.

2-G-38 Does endogenous variation in stress modulate risk and time preferences?

Evgeniya Lukinova[1], Jeffrey Erlich[1]
[1]NYU Shanghai

Objective: It has been argued that one dimension of the cycle of poverty is that poverty is a state of chronic stress. Poor decisions, made under chronic stress, might include carrying high-interest loans, failure to buy health insurance, gambling or drug use. As such, these decisions can contribute to the cycle of poverty. More specifically, a few studies suggest that increased stress may lead to more risk aversion and steeper delay discounting. While the deleterious effects of chronic stress on brain function are well established, much less is known about how chronic stress influences financial decision making. Methods: Here, in a longitudinal design within six weeks period we aimed to incorporate biological mechanisms to improve our understanding of how stress influences economic decisions. We used a combination of decision-making tasks, stress questionnaires, saliva and hair samples within-subject (N=41). We assessed time and risk preferences using hierarchical Bayesian techniques to both pool data and allow heterogeneity in decision making and compared those to cortisol levels and self-reported stress. Results: We found only weak links between endogenous variation in stress and model-based estimates of risk and time preferences. In particular, we found that stress was correlated more strongly with risk preferences, rather than time preferences: cortisol levels in the month preceding the risk task correlated with increased risk tolerance. Also, we found model-free task measures in the short delay task to be moderately related to both hair and saliva cortisol, as well as the stressful life events questionnaire measure. For example, we observed that the stress level one month before the task or the life change units were negatively correlated with the proportion of later choices. Finally, similarly to previous results about subjects' earnings, we established that the curvilinear relationship was preferred to the linear one for waiting time: when cortisol level increased slightly, people waited longer in seconds but when stress increased to higher levels, people waited less. Conclusions: Our results do not directly support the hypotheses that increased stress leads to steeper discounting and increased risk aversion. Instead, we found that endogenous subclinical variation in cortisol only weakly correlates with economic preferences. Moreover, the strongest link we found was that increased cortisol levels correlated with decreased risk aversion, the opposite of our prediction from the literature. Additional multi-task and multi-stress studies with higher N are required to resolve the discrepancies between these results and existing literature.

2-G-52 Structural coherence of the medial forebrain bundle is associated with impulsivity and alcohol consumption in first-year college students

Kelly MacNiven[1], Sarah Hudson[2], Tara Srirangarajan[3], Loreen Tisdall[4], Brian Knutson[3]
[1]Stanford, [2]University of California, Davis, [3]Stanford University, [4]University of Basel

The first year of college is a time of behavioral change for many young adults due to new-found autonomy over choices about what to consume and how to spend time. Previously, we found that structural properties of a specific white matter tract connecting the dopaminergic midbrain and Nucleus Accumbens (NAcc), known as the Medial Forebrain Bundle (MFB), is negatively associated with impulsivity as well as problematic stimulant drug use. Based on these findings, we tested whether MFB tract coherence is related to impulsivity and alcohol consumption in first-year college students (n=59, 28 female, 18-19 years old). To identify the MFB and other control fiber tracts, we applied probabilistic tractography using constrained spherical deconvolution to Diffusion-Weighted Imaging (DWI) data. These methods allowed us to identify bilateral tracts of interest in every subject. Inverse Radial Diffusivity (1-RD) was used as our main measure of tract coherence given its proposed link to axon myelination. Impulsivity was assessed with the Barratt Impulsiveness Scale (BIS), and past 30-day alcohol consumption was assessed via interview using the timeline follow-back method at two timepoints: first on the same day as the DWI scan (baseline), and again four months later (future). Alcohol consumption analyses were performed using overdispersed Poisson regression, after excluding participants who reported never consuming an alcoholic drink (over their lifetime). All analyses controlled for gender. Summary statistics were derived and corrected for multiple comparisons using a permutation-based approach. As predicted, we replicated the negative association of bilateral MFB tract coherence (1-RD) with trait impulsivity ($b=-4.32$, $SE=1.58$, $t=-2.73$, $p<0.05$, corrected). MFB tract coherence was also negatively associated with baseline alcohol drinking (odds ratio (OR)=2.66-6, $t=-2.86$, $p<0.05$, corrected), but not future drinking (OR=0.01, $t=-0.49$, $p=0.62$). Including impulsivity scores in a model did not diminish the association of MFB tract coherence with baseline alcohol drinking (OR=1.83-6, $t=-2.61$, $p<0.05$, corrected), suggesting that trait impulsivity could not account for the negative association of MFB tract coherence with baseline drinking. These findings add to a growing body of work linking structural properties of the MFB to impulsive choice and related behaviors. Future research may build on these findings by exploring how MFB structure supports reward-related function and whether these associations predict or change with impulse-control disorders related to addiction.

2-H-53 The influence of presentation order on learning task structure

Ryan Gallagher[1], William Alexander[1]
[1]Florida Atlantic University

Behavior required in daily activities relies on the ability of an individual to identify structured relationships. For example, a yield sign may initiate stopping or proceeding dependent on the presence of other vehicles: the yield sign serves as context, which prompts the search for more additional concrete information before a response is generated. Research in cognitive neuroscience has demonstrated that the prefrontal cortex serves an imperative role during the integration of information in order to accurately determine appropriate action. While much research has focused on the neural bases supporting this process, less is known about the temporal dynamics that govern it. The purpose of this study is to investigate how individuals assimilate multiple sources of structured information. Participants (N = 40) were instructed to learn through trial and error which on out of three possible keys

was associated with serially displayed pairs of images. Each image sequence contained one abstract and one concrete cue. We defined concrete cues as containing information that more directly minimized response uncertainty compared to abstract cues, which could only minimize uncertainty in conjunction with the concrete cue. We manipulated the presentation order of abstract and concrete informational cues. Subject data was analyzed using traditional statistical tests along with fits of a reinforcement learning model. We found that subjects responded faster ($t(39) = -3.9233$, $p = 0.00034$) and less accurately ($t(39) = -3.2896$, $p = 0.0021$) in the concrete first condition. Fits from our reinforcement model illustrated subjects learned more slowly ($t(39) = -2.8579$, $p = 0.0068$ and responded with a higher degree of noise ($t(39) = -3.0538$, $p = 0.0041$) when concrete cues were presented prior to abstract cues. These results begin to address some of the deficits in understanding regarding the temporal dynamics of information integration and lay a basic foundation by which further studies can be built upon.

2-I-37 Escaping Arrow's Impossibility by Interpersonal Comparison of Neural Utility

Kaosu Matsumori[1], Kazuki Iijima[1], Yukihito Yomogida[1], Kenji Matsumoto[1]
[1]Tamagawa University

Aggregating welfare across individuals to reach collective decisions is one of the most fundamental problems in our society. Interpersonal comparison of utility is pivotal and inevitable for welfare aggregation, because if each person's utility is not interpersonally comparable, there is no rational aggregation procedure that simultaneously satisfies even some very mild conditions for validity (Arrow's impossibility theorem). However, scientific methods for interpersonal comparison of utility have thus far not been available. Here, we have developed a method for interpersonal comparison of utility based on brain signals, by measuring the neural activity of participants performing gambling tasks. We found that activity in the medial frontal region was correlated with changes in expected utility, and that, for the same amount of money, the activity evoked was larger for participants with lower household incomes than for those with higher household incomes. Furthermore, we found that the ratio of neural signals from lower-income participants to those of higher-income participants coincided with estimates of their psychological pleasure by ?impartial spectators?, i.e. disinterested third-party participants satisfying specific conditions. Finally, we derived a decision rule based on aggregated welfare from our experimental data, and confirmed that it was applicable to a distribution problem. These findings suggest that our proposed method for interpersonal comparison of utility enables scientifically reasonable welfare aggregation by escaping from Arrow's impossibility and has implications for the fair distribution of economic goods. Our method can be further applied for evidence-based policy making in nations that use cost-benefit analyses or optimal taxation theory for policy evaluation.

2-I-39 Malleability of human altruism across choice contexts due to social cues

Lisa Bas[1], John-Dennis Parsons[1], Anita Tusche[1]
[1]Queen's University

Introduction. The willingness to act altruistically varies across contexts. What mechanisms drive decisions to bear substantial costs to benefit others in some settings but not in others? Here, we combine an altruistic choice task with formal modeling approaches to examine the role of social cues across contexts. We hypothesize that cues of others' emotional distress - signaling need - may increase generosity. We also investigate differences in people's sensitivity to social factors. Methods. Subjects (N=58) performed an online version of the dictator game. On every trial, subjects chose between two

offers that affected their own and another person's payoff. A generous choice was defined by accepting the offer that benefitted the other at a cost to oneself or by rejecting the offer benefitting oneself at a cost to another. Subjects indicated their preference using a computer mouse. Continuous measures of the cursor position offer a window into the moment-by-moment construction and temporal dynamics of the choice process. To examine the role of social contextual factors, subjects performed the task under two conditions: 'distress' (2 s exposure to dynamic facial expressions of sadness before each trial) or 'control' (no social facial cues) (2 blocks each, 132 trials). Subjects also completed an emotion recognition task and self-report measures of trait differences in mentalizing. Results. Exposure to others' distress increased the generosity compared to the control condition. However, we observed substantial individual differences in the impact of social cues on choices. Subjects with a tendency to spontaneously take others' point of view (better mentalizers) were more affected by social cues during altruistic choice. They were also naturally slower at integrating other's benefits into their decision, possibly due to sufficient leeway for change in the processing speed of others' gains across contexts. Subjects who were responsive to social cues in the dictator game also reported higher real-world donations. This finding suggests that differential sensitivities to social factors may extend to real-world contexts. Notably, emotion recognition performance (d') was comparable for more or less responsive individuals. Thus, our results are unlikely due to differential abilities to correctly identify others' emotional states. Conclusion. The results provide insights into the mechanisms driving variance in altruism across contexts. Our work also points to characteristics of the decision-maker that modify the impact of contextual factors on social behavior. Together, our findings can help to advance research on the malleability of altruism.

2-I-54 Pupil size reflects computational complexity of decisions in humans

Elizabeth Bowman[1], Kristian Rotaru[2], Pablo Franco[1], Carsten Murawski[1]
[1]The University of Melbourne, [2]Monash University

This project examined how computational complexity and cognitive load affect pupil response during complex decisions. 73 participants aged between 18 and 35 years (mean age 22.7 years, 52 female, 21 male) completed 72 trials of the decision knapsack task. Before each trial, participants were required to memorise either a 1-digit or a 6-digit number displayed for 5 seconds. Each trial consisted of a 5- to 7-second grey fixation screen, followed by the knapsack decision task where participants were given up to 25 seconds to decide if a combination of any of the six items displayed on the screen could be found to satisfy a both minimum value and maximum weight constraint. After the knapsack decision screen, participants were required to enter the number they were asked to memorise at the start of the trial. Participants were paid according to the number of trials in which they both correctly recalled the memorised number and made the correct knapsack decision. Knapsack task instances varied in computational complexity (as measured by a control parameter defined by the ratio of normalised capacity and normalised profit), but the number of satisfiable and unsatisfiable instances was the same. Gaze and pupil recordings were performed during the task, sampling at 300Hz with Tobii Pro Spectrum video eye trackers. Pupil recordings were de-blinked, interpolated, smoothed, and downsampled into 2000ms bins for modelling using Growth Curve Analysis (GCA). Mixed-effects modelling of participant responses found a strong effect of instance complexity, and of satisfiability, on trial performance. However, only an inconsistent effect of memory load was found. Participants also spent significantly less time solving knapsack trials of lower instance complexity, and the relationship between instance complexity and time taken significantly predicted performance on a trial. Pupil diameter change was correlated with instance complexity. This difference in pupil response to trials of different instance

complexity was detectable up to 15 seconds before the decision point at the end of the (maximum 25-second) trial, and was greatest in incorrect trials. These results demonstrate that the computational complexity of a task is detected by the participant, and reflected in the effort-related pupil response, well before the decision is signalled through behaviour.

2-I-55 If Artificial Intelligence Kills Consumers?: Decoding the Hot-Cold Organizational Dilemmas from the Multivariate Neural Patterns

Jin Ho Yun[1], Eun-Ju Lee[1]
[1]Sungkyunkwan University (SKKU)

Artificial intelligence (AI) has proven to offer tremendous benefits to many consumers and a nation at large. However, what if AI could kill consumers. In many organizational cases, leaders often face organizational dilemmas that one has to transgress their moral values to acquire financial advantages in favor of the company. This trade-off consists of weighing certain motives to sustain moral principles against the opportunity of receiving the investment offer. Yet, until now how these underlying neuropsychological processes shape leaders' individual moral decisions have remained unclear. Thus, we examined how human brains can decode the Hot-Cold Organizational Dilemmas. We developed the novel Hot-Cold Organizational Dilemmas, in which participants evaluate particular decisions on whether to accept the investment offer at the cost of causing consumers' deaths (Hot Dilemmas: HD) or insignificant discomforts (Cold Dilemmas: CD) by AI technologies. In the functional magnetic resonance imaging (fMRI) scanner, twenty-four (10 women; mean age = 25) participants' neural patterns that dissociate these two types of dilemmas were identified. Participants who imagined themselves as startup leaders read a total of forty HD-CD scenarios (twenty for each dilemma type; varied regarding types of AI technologies ranging from self-driving cars, drones, medical AI, judge AI, robots, etc.). They were also offered investment if decided to release such technologies, otherwise, their startup had to request bankruptcy if one decided not to release this technology. Subsequently, they made judgments about the investment amount to accept. In this context, leaders face the conflict between the temptation of accepting the investment offer for preventing bankruptcy (i.e., utilitarian judgment) and not accepting the offer for preventing consumers' death or discomfort (i.e., deontological judgment). Employing machine learning methods, we found that the decoding accuracy of neural patterns pertained to HD-CD decisions correlated significantly with differences in not accepting the investment offer between dilemma types in brain areas associated with cognitive control and theory-of-mind (ToM), such as the temporoparietal junction and temporal pole. Our findings suggest that cognitive control and ToM processes are both needed to make utilitarian judgments in CD and to make deontological judgments in HD. Furthermore, individual differences in deontological preferences differentially modulated response-specific signals in the medial prefrontal cortex according to those who did not accept the investment offer.

2-I-56 Two distinct patterns of EEG activity emerge during value-based choice, neither related to evidence accumulation

Romy Froemer[1], Matthew Nassar[1], Benedikt Ehinger[2], Amitai Shenhav[1]
[1]Brown University, [2]University of Stuttgart

Objective Research into value-based choice and its neural correlates often centers on a unitary process subsuming the accumulation of evidence about one's options until a choice is made. The centro-parietal

positivity (CPP) is one putative index to this evidence accumulation process. At odds with a unitary view, our previous fMRI research suggests two distinct forms of evaluation during value-based choice: (1) overall appraisal of one's options, and (2) comparing between those options for the current choice. Here we test whether these processes are also temporally dissociable, with appraisal-related processes tied to the time one's options appear and choice-related processes tied to the time a decision is made. In light of this dissociation, we reevaluate the CPP as an index of evidence accumulation. Method We recorded EEG while participants (N = 39) made choices between previously individually rated consumer goods. Afterward, participants rated their overall appraisal (liking) of each set, how anxious they felt making the choice, how confident they were in this choice. Using PCA, we reduced the dimensionality of measures to two components, broadly indexing variance associated with the appraisal of items (e.g. overall value and set appraisal) and choice between items (e.g. value difference and confidence). We used those components as regressors for stimulus- and response-locked EEG activity. Results As predicted, activity locked to the onset of the stimuli correlated with our appraisal-related PC, and activity locked to (and preceding) the response with our choice-related PC. While neither activity patterns were consistent with canonical indices of evidence accumulation, we show that typical CPP evidence accumulation patterns can artificially emerge as stimulus and response locked effects occur closely in time. In a separate dataset we further show that observed signatures of evidence accumulation vanish when accounting for such overlap through joint modeling of stimulus and response-locked activity. Conclusion Our results provide further evidence that neural activity during value-based choice reflects more than a unitary choice process. Neither of the spatiotemporally distinct patterns associated with appraisal and choice conformed with canonical signatures of evidence accumulation. They might instead reflect affective and metacognitive appraisals of one's options and (latent) choice. Our findings call for caution when interpreting neural correlates of choice value, a critical re-examination of putative EEG evidence accumulation signals, and for taking a broader, more integrative perspective on value-based choice.

2-I-57 Cognitive Effort Discounting in ADHD-Diagnosed and Healthy Control Adolescents

Deborah Sevigny-Resetco[1], Suzanne Mitchell[1]
[1]Oregon Health and Science University

Objective: Differences in reward valuation and processing have been linked to an ADHD-diagnosis. However, subjective preferences and biases that may underlie these differences remain unclear. Previous studies with this population have focused on rewards earned after completing tasks that require physical effort, rather than cognitive effort. Studying the latter, however, may provide more robust insights into the observed differences in cognition between ADHD and controls, including differences in attention, working memory, inhibition and cognitive flexibility, and self-regulatory processes. Our study aimed to address this knowledge gap by examining cognitive effort discounting: choices between minimal effort rewards and larger rewards with higher effort costs. Methods: This study recruited ADHD-diagnosed and healthy control participants aged 16-21 years. Participants performed and rated six variants of a Sustained Attention Task and a Working Memory Task, each with varying difficulty levels, produced by varying display rates (inter-stimuli-interval: 750 or 1000 ms and stimulus durations: 200, 400, 800 ms). A variant of each task, rated as moderately effortful, was identified for each participant and then practiced. Participants imagined performing each task for 1, 5, 10, and 20 minutes and rated the imagined experience. Participants subsequently completed two Cognitive Effort Discounting Tasks examining preferences between large monetary rewards ($10, $25,

$50) requiring task performance for a specific duration (1, 5, 10, 20 minutes) and smaller monetary rewards obtained with no effort required. Results: As expected, analyses demonstrated that variants with more rapid display rates were perceived and rated as more difficult by both groups on both tasks. There were no significant differences between ratings of task liking, discomfort, mental effort, nor confidence in the ability to perform either task between groups. Systematic cognitive effort discounting curves were obtained for both groups, but only a trend toward increased discounting of effortful rewards was found in the ADHD group. Analyses of indifference points revealed ADHD-diagnosed individuals discounted more steeply for the larger rewards ($25 and $50) but not when the rewards were relatively small ($10). Conclusions: These data suggest a difference in discounting, despite equating tasks on subjective effort. Future work will examine additional correlates (eye-tracking, heart rate, and skin conductance) and personality measures to isolate possible factors that may contribute to these differences in decision-making.

2-I-58 Decoding monetary and social reward anticipation using whole-brain multivariate pattern analysis

Huan Wang[1], Elizabeth Blevins[1], Jeanne Tsai[1], Brian Knutson[1]
[1]Stanford University

Recent meta-analyses comparing monetary and social reward anticipation based on the Monetary Incentive Delay (MID) task and the Social Incentive Delay (SID) task with Functional Magnetic Resonance Imaging (FMRI) reported that anticipation of both reward types engaged a domain-general reward anticipation network (Gu et al., 2019). Nevertheless, some studies reported domain-specific activation of social cognitive regions during social reward anticipation (e.g., Spreckelmeyer et al., 2013). Given this inconsistency, we applied the Support Vector Machine (SVM) with recursive feature elimination (De Martino et al., 2008) to the whole-brain FMRI data on optimized versions of the MID and SID tasks to address whether social cognitive regions will be implicated in social but not monetary reward anticipation. Fifty-eight healthy adults participated in the study. We trained SVM classifiers for each task, and tested generalization performance using leave-one-out cross-validation across subjects. We submitted preprocessed whole-brain FMRI data from the first and second volume acquisitions of the anticipation phase to SVM in order to assess potential temporal changes in features associated with reward anticipation. Additionally, we compared the univariate with multivariate results by extracting information from predicted Volumes Of Interests (VOIs) - Nucleus Accumbens (NAcc), associated with general reward anticipation, and the temporoparietal junction (TPJ), associated with social cognition. Classifiers for both the MID and SID task discriminated high versus no reward anticipation better than chance (50 percent) for the first (MID: M = .79, SID: M = .70; ps<.001) and the second acquisition volumes (ps<.001). Furthermore, VOI analyses revealed that TPJ was involved for social but not monetary reward anticipation, although only based on multivariate results (ps<.03) but not univariate results. By contrast, both univariate and multivariate results indicated that the NAcc was similarly implicated for monetary and social reward anticipation (ps<.003). The current results provide partial support for the domain-general reward anticipation network view, but also suggest some specificity. Univariate VOI results suggest similar activation of NAcc for monetary and social reward anticipation, and no difference in the TPJ. Multivariate analyses, however, suggest that TPJ is involved only in classifying social reward anticipation. These findings shed light on both domain generality and specificity in the context of reward anticipation, and highlight the importance of applying multivariate techniques to FMRI data to reveal potential hidden information.

2-I-61 Meta-Analysis of Explore-Exploit Decisions Reveals Convergence in the Salience Network

Daniel Sazhin[1], Amanda Nguyen[1], David Smith[1]
[1]Temple University

Introduction: Explore-exploit tasks studied through n-armed bandits (Daw et al., 2006) and foraging tasks (Stephens and Krebs, 1986) do not have apparent behavioral convergent validity (von Helversen et al., 2018), and it is unclear what leads to these divergent results. Across explore-exploit studies there seems to be consistent brain activation, yet it remains unclear whether there is convergence across studies. To address this issue, we conducted a coordinate-base meta-analysis across neuroimaging studies using n-armed bandit tasks and foraging tasks. Methods: In the current study, we used the PRISMA system (Liberati et al., 2009) to conduct an extensive literature review of n-armed bandits and foraging tasks. We used the following search terms on PubMed: (fMRI OR functional magnetic resonance imaging OR neuroimaging OR brain OR neural) AND (n-armed OR exploration-exploitation OR explore-exploit OR multi-armed OR foraging). This search identified 1956 publications. After exclusions were applied, 14 studies passed through the eligibility requirements. Next, we coded these studies as explore-exploit (N = 9) and foraging (N = 5). Analyses were performed using the Seed-based d Mapping (SDM) software. Results: Given the relatively small number of foraging studies, our preliminary analyses did not include these coordinates. We found that explore-exploit studies were associated with convergence in the ventral striatum, anterior cingulate cortex (ACC) and the intraparietal sulcus. Next, we conducted subgroup analyses to examine the specificity of these effects. In the explore > exploit subgroup (N = 9), we found convergence in activation in the right anterior insula (AI), dACC, and the right precuneus. In the exploit > explore subgroup (N = 4), our results indicate convergence in activation in the left precuneus and the right AI. Conclusions: The results were somewhat contrary to our hypotheses as areas associated with executive function such as the vmPFC and dlPFC were not found in our analysis. However, the AI and ACC feature prominently in the salience network (Seeley, 2019) and were significant in both explore and exploit conditions. Further analyses are required to assess conjunction and convergence between conditions. in our explore and exploit analyses respectively Overall, these results suggest that the salience network may help distinguish differences in explore-exploit decisions.

2-I-62 Willingness-to-pay for social experiences: How social cognitive functioning relates to individuals' choices between social and non-social experiences

Melanie Ruiz[1], Joanne Stasiak[2], Peter Sokol-Hessner[3], David Smith[4], Dominic Fareri[1]
[1]Adelphi University, [2]University of California-- Santa Barbara, [3]University of Denver, [4]Temple University

Objective: Human behavior is largely driven by a desire for social connection. Navigating the social world in pursuit of positive social experiences requires engaging in social valuation--i.e., assigning differential value to socially relevant stimuli. Difficulty engaging in social valuation is associated with negative psychological and physical health outcomes, but the degree to which people vary in their valuation of social relative to non-social experiences has not been directly quantified. Critically, the relationship between available social resources (i.e., social support, social abilities) and computational processes supporting social valuation are surprisingly understudied. Here, we investigated the degree to which people would spend earned money on social relative to non-social experiences, and whether this varied with individual differences in social resources. Methods: Participants (N=101, 47 female, ages 18-35) recruited from Adelphi University and Prolific.co engaged in a willingness-to-pay (WTP) task to assess

the relative value of social to non-social experiences after first earning an endowment ($6.00). On each of 50 trials in the WTP task, participants chose between one of two hypothetical experiences, one social (e.g. going to dinner with friends) and one non-social (e.g., watching TV at home alone), each associated with a different price (range: $.01-.05) that would be deducted from their endowment. We also assessed individual differences in self-reported social functioning (e.g., autistic traits) and perceived social support. Results: On average, participants spent more of their endowment on social (M = 0.029) relative to non-social (M = 0.027) experiences ($t(100)=7.732$, $p < .001$, $d = 0.43$) and favored social experiences significantly more than chance ($t(100)=2.304$, $p= 0.023$, $d =0.46$). Additionally, we observed a positive correlation between levels of perceived social support and the likelihood of choosing social over non-social experiences ($r(100) = .195$, $p = .05$). Finally, those endorsing autistic traits to a higher degree were less likely to choose social over non-social experiences ($r(100) = -.40$, $p < .0001$). Conclusions: Taken together, our preliminary results suggest that social resources may play a significant role in the tendency to place value on social experiences. These findings have implications for understanding altered social valuation processes in conditions associated with social difficulties (i.e., autism, depression). Future analyses will compute the point of subjective equivalence for social and non-social experiences and will examine the role of loneliness and mentalizing abilities on social valuation.

2-I-63 Task complexity and experience dictate the use of online, versus offline, planning in humans

Jungsun Yoo[1], Aaron Bornstein[1]
[1]University of California, Irvine

Objective: How much do our decisions depend on deliberation in the moment, as opposed to reasoning performed ahead of time? In reinforcement learning (RL) theory, these approaches are distinguished as online and offline planning (Sutton & Barto 2018). Understanding when each kind of planning is used can help predict how decisions may change in response to e.g., time pressures or context effects. Though recent work suggests that artificial agents (trained across millions of examples of a given task) are overwhelmingly guided by offline planning (Hamrick et al 2020), an open question is whether and when humans (with fewer direct examples to learn from) use online planning. Method: To approach this question, we developed a novel measure of online planning and examined how it differed with task complexity and across experience within a given task. We reasoned that online planning would be more prevalent when the task is more complex (operationalized as more possible states), and as internal representations developed with learning. Four variants of the standard 'two-step' task (Daw et al 2011; N=100 each) were conducted. Variants differed in the number of possible first-stage states - 1, 3, 6, or 10 - and thus the degree to which participants were able to anticipate the upcoming decision problem before the trial began. We computed both the original measure (w), which combines online and offline planning, and a new, deliberation-sensitive measure, the online planning index (OPI). Results: Consistent with our predictions, we observed that participants were more likely to use online planning as a function of task complexity and within-task experience. First, while the the tendency to use planning of any kind was directly related to performance in the canonical 1-state version of the task (direct effect w → score: $\beta = 20.81$, $p < .001$), in more complex variants this relationship grew progressively more mediated by the degree of online planning, to the extent that the 3- and 6-state tasks showed no direct effect from w to performance but a significant effect from w to online planning (3: $\beta = 0.25$, $p < .001$; 6: $\beta = 0.37$, $p < .001$, respectively), and a mediation effect was observed in the 10-state task (w → OPI: $\beta = 0.23$, $p < .001$; OPI → score: $\beta = 14.75$, $p = .01$). Second, in the 10-state task, the mediation of online planning evolved

quadratically with experience. Conclusions: These results suggest that humans learn to use online planning in complex environments. Our findings have implications for the generalizability of decision-making models to naturalistic settings, and for the correspondence between decision-making in humans and artificial agents.

2-I-65 Humans learn prototype states and dynamically generate decision boundaries during economic choices

Warren Pettine[1], Dhruva Raman[2], A Redish[3], John Murray[1]
[1]Yale School of Medicine, [2]Cambridge University, [3]University of Minnesota

Objectives Decisions in cue-rich environments require an inferential state recognition process. Psychological latent-state learning models typically compare incoming sensory and context information against idealized prototypes, assigning identity to the closest. In contrast, many in the machine learning community have embraced a decision-boundary approach where the weight of an attribute is determined by its usefulness in state-discrimination. A third model, derived from category learning studies, compares incoming sensory information to multiple state exemplars and selects the state most closely matching them. We designed three experiments to differentiate these hypotheses and deployed them to human subjects online. Experimental observations are best captured by a model that forms prototype states and uses them to dynamically generate decision-boundaries during decisions. Methods Model: Our latent-state reinforcement learning model forms prototype states to maximize reward received across trials. Each trial, the model develops a candidate list of potential states, then computes the mutual information across states for each attribute to form decision-boundary weights. Exp1 (n=63): Subjects categorized stimuli that differed along two attributes (color, shape), then generalized categories to stimuli with a novel texture. A purely prototype model predicts difficulty with generalization, while a decision-boundary model does not. Exp 2 (n=56): In Block 1, subjects categorized stimuli according to their shape. In Block 2, subjects categorized a new set of stimuli according to color. In Block 3, subjects made novel comparisons among all categories observed in the prior blocks. A purely discriminative model predicts confusion between states learned separately in each block when the only difference is a previously non-informative attribute (texture). A prototype model generalizes without issue. Exp 3 (n=70): We modified Exp 2, adding a new attribute (size), and arranged stimuli such that an exemplar model would introduce specific biases in the errors. Results Exp 1: Human subjects demonstrated the use of discriminative decision-boundaries during generalization. Exp 2: Subjects generalized through prototype states rather than decision-boundaries. Exp 3: Subjects formed prototype states, rather than exemplars. Our computational model effectively performed all experiments and qualitatively captured human behavior. Conclusions Human generalization requires prototype states with the formation of discriminative decision boundaries. Our work has key relevance for how subjects form concepts of economic goods, and what attributes are utilized in decisions.

1-I-96 The consequences of base rate neglect on sequential belief updating and real-world beliefs

Brandon Ashinoff[1], Justin Buck[1], Michael Woodford[1], Guillermo Horga[1]
[1]Columbia University

Study Objective: Base-rate neglect (BRN) is a pervasive bias where prior knowledge is under-incorporated into newly formed beliefs. Surprisingly, the underlying neurocomputational mechanisms of BRN remain poorly understood, as are its effects on sequential belief updating and real-world beliefs. A

Bayesian formalization of BRN predicts that belief trajectories should depend on the order that evidence is presented; this should manifest as a recency bias defined by greater weighting of newer evidence and lingering uncertainty over time due to a form of prior-dependent updating, suggesting that base-rate neglecters may be prone to forming unusual long-term beliefs in the real world. Methods: We developed a novel 'urn-and-beads' task that systematically manipulated the evidence order of 8-sample sequences and used an incentive-compatible belief-elicitation procedure. Results: Using model-based and -agnostic measures, two independent online studies (n=151; n=116, reported) confirm the prediction that BRN is expressed as a recency bias (p=8.0369x10-5; sign-rank) and as prior-dependent updating (p=2.7286x10-4; mixed-effects). While individuals exhibited BRN on average (p=7.3592x10-4; sign-rank), there was interindividual variability that scaled with the magnitude of the recency bias (ρ=-0.53, p=8.8137x10-8; spearman) and prior-dependent updating (ρ=0.24, p=0.02; spearman). Since BRN induces inaccurate beliefs, it is commonly framed as a suboptimal process. However, bounded-rationality models frame it as an optimal response to limited internal resources. Consistent with this, a noisy-sampling model whereby BRN results from rational discounting of noisy internal representations of prior probabilities provided an accurate description of our combined dataset (n=267). Specifically, variability in probability estimates correlated with BRN consistent with the noisy sampling model (ρ=0.23, p=1.2215x10-4; partial spearman). Thus, BRN may be driven by imprecision in the internal representation of prior knowledge. We also examined self-reported propensity for unusual beliefs outside the laboratory and found individuals (n = 116) with more BRN (ρ=-0.21, p=0.03) and noisier prior representation (ρ =-0.27, p=0.005; partial spearman) tended to report more unusual beliefs in their daily lives. Conclusions: Our results demonstrate a behavioral signature of BRN in sequential belief updating that conforms to Bayesian predictions and that this bias has a measurable influence on real-world beliefs. Further, our results suggest a candidate mechanism where BRN reflects a rational response to imprecise internal priors.

2-K-59 Evaluating the evidence for preference-based attentional capture in binary choice

Xiaozhi Yang[1], Blair Shevlin[1], Ian Krajbich[1]
[1]The Ohio State University

What is the relationship between gaze, value, and the decision process? Some researchers have argued that gaze amplifies value to drive the decision process, while others argue that gaze is drawn to options that are more likely to be chosen. In a recent article, Westbrook et al. (2020, Science) argued for a hybrid two-stage choice process, with an early stage where gaze amplifies information and a late stage where gaze is directed towards the to-be-chosen option. However, their analyses were not sufficient to support this claim because they would identify two stages even in one-stage data. Here, we took a more general approach for detecting whether gaze is drawn to preferred options (Sepulveda et al. 2020, eLife) and applied it to 11 datasets. Ultimately, we found support for preference-driven gaze in only a few specific cases, indicating that this is not a general feature of value-based decision-making. Westbrook et al. used a bifurcation analysis to identify two stages of the decision process, finding that gaze amplification is strongest pre-bifurcation, while gaze bias is strongest post-bifurcation. However, we find that the model parameters from this approach are highly sensitive to the selection of the bifurcation point and may identify such differences even from simulations of a single-stage model. Thus, we argue that this approach cannot be relied on to distinguish between cognitive mechanisms. Sepulveda et al. used an alternative analysis, examining how the correlation between gaze location and value difference evolves over the course of a trial. In their gaze-contingent paradigm, they find that this correlation

increases over time. Applying the same analysis to the Westbrook dataset and simulations of a two-stage model, we also found gaze-value correlations that increase over time. However, when we applied the same analysis to other datasets we found little additional evidence for this phenomenon. The datasets that we examined come from binary-choice eye-tracking experiments, some with options broken down into distinct attributes. These include choices from consumer goods, social allocations, risky prospects, and conditioned stimuli. We found that most datasets do not display any correlation between gaze and value. The only datasets that do contain such patterns are those where subjects may have employed an alternative, search-based choice strategy. In summary, our results suggest that goal framing and task structure can impact the relationship between gaze and choice, and that preference-based attentional capture only occurs in certain settings.

2-L-60 How multi-slice FMRI acquisition can compromise detection of mesolimbic reward responses

Leili Mortazavi[1], Tara Srirangarajan[1], Brian Knutson[1]
[1]Stanford University

Simultaneous multi-slice acquisition allows researchers to collect Functional Magnetic Resonance Imaging (FMRI) data with higher temporal resolution and faster speed of data collection, as compared to traditional single-slice protocols. As a result, multi-slice acquisition has become increasingly widespread for neuroimaging studies in recent years. However, this innovation may come at a cost to signal quality in subcortical regions of the brain. In a meta-analysis (N=44 studies; cumulative n=5005 subjects) and direct comparison of FMRI studies using the Monetary Incentive Delay (MID) task, we previously found that relative to conventional single-slice acquisition, multi-slice acquisition was associated with significantly decreased effect sizes for reward anticipation in the Nucleus Accumbens (NAcc) as well as for reward outcomes in the Medial PreFrontal Cortex (MPFC) by more than half (SNE 2020). To try to isolate the source of the decrement, and control for variables related to task, magnet, and subjects, we conducted a pre-registered within-subject and within-session comparison of single-slice and multi-slice data (multiband factors=4 and 8; n=12 subjects), using a well-characterized and reliable reward task (i.e., MID Task; Knutson et al., 2000). Results indicated that acquisition of multi-slice data diminished reward versus neutral anticipation effect sizes in mesolimbic projection areas including the NAcc, and that the association between single versus multi-slice acquisition and reward anticipation effect size was fully statistically mediated by reductions in temporal signal-to-noise ratio. Together, these findings suggest that researchers who wish to probe mesolimbic activity with FMRI should opt for single-slice rather than multi-slice acquisition protocols. This growing literature also has broader implications for inferring mesolimbic activity during related tasks and rest, for summarizing the past literature, and for using FMRI to track individual differences in reward-related brain activity.

2-L-66 Is Poverty Low Control over Queue Priorities?

Marcos Gallo[1], Kavya Rajagopalan[1], Colin Camerer[1]
[1]California Institute of Technology

The poor face multiple uncontrolled demands on their time and money. Recent work has hypothesized that resource scarcity creates biases in decision making, increasing attention towards scarce resources at the cost of attention for unrelated, but important, aspects. We provide an updated version of this theory, and we argue that these changes in the decision process are a result of low control over priority queues. Those with more resources are faster at completing some tasks and are able to delegate to

others. We develop a model and test it with a novel task suggesting that distortions in the perception of urgency and importance deviate from optimal behavior.

Poster Session #3

3-A-72 Can brain activity of financial professionals forecast stock market performance?

Leo van Brussel[1], Ale Smidts[1], Roeland Dietvorst[2], Maarten Boksem[1]
[1]Erasmus University Rotterdam, [2]NN Investment Partners

Objective: One of the major aspirations of investors is to better forecast stock performance. Interestingly, the emerging field of neuroforecasting shows that neural predictors of choice can scale up to explain market-level behavior, above and beyond traditional measures (Genevsky & Knutson, 2018). However, whether neural responses can be identified to forecast stock market performance remains to be investigated. In this study we set out to test whether neural responses to real-world investment opportunities are indeed related to their performance on the stock market. Methods: We invited 36 financial professionals to participate in a Stock Performance Task (SPT) while undergoing fMRI. The SPT consisted of 40 anonymized, real-world investment cases of companies either overperforming or underperforming in their market segment. Per investment case, five consecutive information screens were presented: (1) Company profile, (2) Stock market price graph, (3) Financial fundamentals, (4) Relative valuation and a (5) News item. Next, participants were asked to predict whether the stock would be "overperforming" or "underperforming" in its market segment in the next year. After the experimental task, participants completed several self-report measures (e.g rational vs. experiential thinking style; Novak & Hoffmann,2009). Results: We report initial findings related to participant's choices. We found that our sample of financial professionals could not predict stock market performance of the investment cases beyond chance. Individual performance in the task showed no correlation with participant's thinking-style or decision confidence. Next, we used multilevel modelling to analyze decisions on a trial-by-trial basis. Following the neuroforecasting literature, we entered two key regions of the valuation system (vmPFC and VS; Bartra et al., 2013) as neural predictors. We found vmPFC activity during the information screens to be related to choice. Particularly, vmPFC activation during the news item screen increased the probability of participants choosing "overperforming". A whole-brain fMRI analysis confirmed involvement of the vmPFC during evaluation of the information screens to be related to participants' choices. Conclusions: Our preliminary results show involvement of the brain's valuation system when financial professionals evaluate information pertaining to real-world investment cases, independent of thinking style. On average, we found that participants were not able to predict stock performance above chance. We are currently investigating whether neural activity in the valuation system of the brain can be related to performance on the stock market.

3-B-94 Consumer responses to narrative advertising: do narrativity level and channel type play a role?

Aline Simonetti[1], Hossein Dini[2], Diego Candia-Rivera[3], Luis Bruni[2], Enrique Bigne[1]
[1]University of Valencia, [2]Aalborg University, [3]University of Pisa

Narratives can be a powerful mechanism for persuasion. When correctly applied to the advertising field, it potentially increases ad and brand evaluations. Virtual reality (VR) used in advertising can foster imagery richness and vividness compared to traditional video ads. Irrespective of the exact mechanism behind ad persuasion, cognitive and affective processes are involved in narrative processing. We aim to

analyze narrative advertising varying in its narrativity level and explore possible channel effects--ads presented on a computer screen (PC) vs using a VR headset. Methodology: A 2 (narrativity level: high--HL vs low--LL) x 2 (channel: PC vs VR) within-subjects (WS) experiment was conducted. 32 participants (13 female, Mage=26.84, SD=4.33) of 16 nationalities watched 12 2D real video ads (six HL, six LL; ~60s each) in PC and VR. A 10-item questionnaire was answered after each ad and galvanic skin response (GSR) was recorded throughout the study. The questionnaire assessed consumers' response to ads, namely: affect, processing, signals, and brand attitude. Results: For the questionnaire metrics, a WS generalized linear mixed model (GzLM) was used (fixed factors: narrativity level, channel, and the interaction term; random factors: person and brand). There was a main effect of narrativity level on (F(1,759), p-value, direction of the effect): Ad Liking (12.30, <.001, HL>LL); Ad Sharing (9.86, <.01, HL>LL); Ad Attention (16.09, <.001, HL>LL); Ad Immersion (9.78, <.01, HL>LL); Ad Acceptance (47.23, <001, HL>LL); Brand Trust (67.63, <.001, HL<LL). There was a main effect of channel on: Ad Immersion (5.57, <.05, VR>PC); Brand Reaction (4.81, <.05, PC>VR). No effects were found for Ad Understanding. Interaction effects were present in four metrics. For the GSR, we used a non-linear estimation of the number of spontaneous fluctuations (PsPM software). The results of the analysis (WS-GzLM without the random factor brand) showed the main effect in the channel, where VR elicited more spontaneous GSR (F(1,518) = 65.30, p<.001). Conclusions: We found an overall better performance of ads with high levels of narrativity on creating more positive ad perceptions, such as liking, intention to share, attention paid, immersive feelings, and acceptance. But those ads performed worse on trust perceptions. Regarding channel, it had a small effect on immersion feelings and brand reactions. Participants felt more immersed into the ads when watching them on VR, but they reported higher favorable reactions when watching them on a PC. Although the GSR revealed that videos watched on VR led to higher arousal, this was not enough to influence ads' conscious perception.

3-C-73 Neurocomputational mechanisms engaged in detecting cooperative and competitive intentions of others

Rémi Philippe[1], Rémi Janet, Koosha Khalvati[2], Rajesh P. Rao[2], Jean-Claude Dreher[1]
[1]CNRS, [2]University of Washington

Humans generally interact without knowing the intentions of others in terms of competition or cooperation. This can make situations difficult to predict, especially when the intentions of others fluctuate over time between cooperative and competitive goals. Little is known about the neurocognitive mechanisms that underlie inferring intentions of others in such cooperative/competitive situations. Here, we used computational modeling with fMRI in a new interactive dyadic game to investigate how the brain dynamically tracks the cooperative/competitive intentions behind the actions of others. Participants were required to select the same card as a co-player from two available cards. Unbeknownst to them, the co-player was an artificial agent that alternated, without signaling, between cooperative and competitive behaviors (matching pennies vs hide and seek). To win, participants had to adapt their strategies appropriately. Comparison between computational models demonstrate that during such inferences, participants arbitrated between competitiveness and cooperativeness of others on a trial-by-trial basis to weight decisions and adjust behavior. Model-based fMRI analyses revealed that the striatum implements an arbitration process between two expert systems dedicated to track competitive vs cooperative strategies. The dorsolateral prefrontal cortex and temporo-parietal junction, regions involved in theory of mind, differentiated the two modes of interactions when receiving feedback. Moreover, functional coupling between the striatum and the dorsolateral prefrontal cortex

increases when engaging in attributing intention to make strategic decision. Together, these findings indicate that the interaction between theory of mind brain network and striatum is key to dynamically monitor the reliability of the predictions of different expert systems in terms of competition and cooperation of others.

3-D-75 Asymmetric social contagion effect in risk but not in ambiguity decision making

Deng Pan[1], Qingtian Mi[1], Lusha Zhu[1], Jian Li[1]
[1]Peking University

Objective People's choice behaviors are profoundly intertwined with observing and learning from others in the social environment. Recent studies have shown that people change their initial risk preferences to better align with the risk attitude of other agents after observing others' risk-related decisions (i.e., risk contagion effect). However, it remains unknown whether such contagion effects are symmetric between observing conservative and adventurous individuals, and, furthermore, whether the contagion effect can be generalized to other types of uncertainty such as ambiguity. Here, by combining behavioral results and psychological modeling, we investigate the critical question of whether observing decisions made by conservative and adventurous people yields asymmetric effects in different decision frames (gain & loss), given the ubiquity of social learning. We also investigate the similarities and differences between the social contagion effect of risk and ambiguity. Methods Two independent behavioral experiments (N1=41, N2=60) were conducted to address these questions (Exp. 1: risk, Exp. 2: ambiguity). In both experiments, participants were instructed to decide between a gamble option and a sure option presented in a particular decision frame. Meanwhile, participants also needed to observe and predict decisions made by another agent. We mainly focus on how observing other people's decisions (conservative/adventurous) in different decision frames (gain/loss) would affect participants' attitudes towards uncertainty (risk/ambiguity). Results We found that individual's attitudes towards risk and ambiguity were shifted after observing other's decisions. Intriguingly, participants were more readily aligned with a risk-averse social partner than a risk-seeking partner both in the gain and loss frame, indicating an asymmetry of social influence from other's risk attitudes. In contrast, we found that participants were equally influenced by observing the decisions made by the ambiguity-averse and ambiguity-seeking social agents in both decision frames. Conclusions These results suggest that social influence changes human attitudes towards risk and ambiguity differently. Specifically, the diversion is manifested by the asymmetric contagion effect on risk but the symmetric effect on ambiguity.

3-D-76 The functional roles of right DLPFC and VMPFC in risk-taking behavior

Aline Dantas[1], Elisabeth Bruggen[1], Peiran Jiao[1], Alexander Sack[1], Teresa Schuhmann[1]
[1]Maastricht University

Many of our daily decisions involve some form of risk, may it be investing in stocks, gambling in a casino or trying a new cuisine. This risk-taking behavior is part of a decision process that has been associated with activity changes in specific prefrontal regions of the brain, including the ventromedial prefrontal (VMPFC) and the dorsolateral prefrontal cortex (DLPFC) (Fecteau et al., 2007; Gianotti et al., 2009), likely activated by the involved sub processes of valuation (DMPFC) and executive control (DLPFC) (Berkman, 2018; Clark et al., 2017; Gallagher et al., 2009) Using non-invasive brain stimulation methods (NIBS), the functional role of the DLPFC in risk-taking behavior consistently revealed that inhibition of the DLPFC leads to an increase in risk taking behavior (Boggio et al., 2010; Fecteau et al., 2007; Gilmore et al.,

2018). In contrast, the specific, potentially different role of the VMPFC for risk taking behavior is less well established. Although many brain imaging studies take the activity of these two prefrontal areas as indicators of different functions during risky choices (Hare et al., 2009; Hutcherson et al., 2012), this differential functional contribution is yet to be confirmed. Here, we used continuous theta burst stimulation (cTBS) to inhibit either the VMPFC or DLPFC during the execution of the computerized Maastricht Gambling Task (MGT - Dantas et al., 2021), allowing us to estimate participants' risk-taking behavior. This enables us to investigate the specific and differential effects of TMS-induced inhibition of the DLPFC and VMPFC on both, risk-taking behavior as a whole, and the valuation process in particular. We hypothesized that, compared to sham stimulation, VMPFC inhibition would lead to a reduction in risk-taking behavior by reducing the participants' choice of average values, whereas the TMS-induced inhibition of the right DLPFC (rDLPFC) would lead to an increase in risk taking due to a reduction in cognitive control (Knoch et al., 2006). In a within-subject design, 33 participants were asked to complete the MGT before and after receiving either VMPFC, rDLPFC or sham stimulation in separate, counterbalanced sessions. Transcranial magnetic stimulation (TMS) was applied with a MagPro X100 stimulator using a double cone coil (Magventure, Denmark) at 100% resting motor threshold. Stimulation of both the rDLPFC and the VMPFC led to an increase in risk-taking behavior (beta = -0.47, t(8847) = -2.98, p < .01beta = -0.45, 95%, t(8847) = -3.01, p < .01; respectively). We also found a significant medium positive increase in the average valuation after both DLPFC (beta = 0.67, t(8847) = 2.63, p < .01) and VMPFC stimulation

3-E-68 Goal-dependent recalibration of hippocampal representations facilitates self-control

Micah Edelson[1], Todd Hare[1]
[1]University of Zurich

Objective: The neurobiological interactions between mnemonic and decision-making processes are important for determining individual behavior. One domain where this may be especially relevant are decisions that require self-control because they rely on the ability to select and control the internal information utilized when making a choice. For example, when dieting and maintaining a goal to eat healthfully, successful self-controllers may be able to prioritize the recall of a subset of experiences in which chocolate cake was less tasty than usual or a salad was particularly delicious. Theoretically, this would result in mnemonic representations of the past and/or future forecasted tastiness that are more similar for chocolate cake and salad (because the tastiness of salad is increased and that of chocolate cake decreased). We refer to this information control process as goal-dependent recalibration and predict that it should occur in the hippocampus given that region's well-established role in episodic memory. Methods: We tested our hypothesis by conducting representational similarity analyses of fMRI activity within the hippocampus while humans (N = 44) performed blocks of 50 food rating or dietary choice trials. Results: We found that people do flexibly adjust their multivariate hippocampal representations of the palatability of food items in a goal-dependent fashion. On average, the entire sample of participants showed significant goal-dependent recalibration of taste representations in the hippocampus between trials in which the goal was to rate the tastiness versus rate the healthiness of a food item (Cohen's d= 0.69, posterior probability = 0.999). Critically, this goal-dependent recalibration was also associated with actual choices in the face of self-control challenges. In fact, participants with high, relative to low, levels of self-control did represent palatable and unpalatable food items more similarly within the hippocampus when facing self-control challenges (Cohen's d = 0.8, posterior probability = 0.989) and the amount of goal dependent recalibration was strongly predictive of self-

control success (out-of-sample accuracy = 69%, probability > chance = 0.996). Conclusions: Our results provide new insights into the processes involved in self-control. They indicate the need to incorporate memory and hippocampal representations into mechanistic frameworks of self-control alongside the longstanding focus on i) regulatory functions supported by the prefrontal cortex and, ii) antagonistic interactions between prefrontal cortex and subcortical brain regions.

3-E-69 Does COVID-related stress affect self-control and the ability to make healthy food choices?

Marie Falkenstein[1], Felix Nitsch[2], Leonie Koban[1], Aiqing Ling[3], Tobias Kalenscher[2], Hilke Plassmann[4]
[1]Sorbonne University, [2]Heinrich-Heine-Universität Düsseldorf, [3]University College Dublin, [4]INSEAD

Objective: Experimentally inducing acute stress has been previously found to impair self-control. Research on the effects of chronic stress on self-control is sparse but research on more general decision-making suggests a negative impact. The body responds to stressors by stimulating the hypothalamic-pituitary-adrenal (HPA) axis releasing cortisol which is commonly known as the stress hormone. The ratio of cortisol to cortisone can be used to assess the activity of the HPA axis. Here, we investigated whether the Covid-19 crisis is an exogenous stressor that impacts self-control. Methods: Our study consisted of two parts: (1)An online study carried out across China and France summer 2020. Participants were asked to make choices between foods and report their chronic stress levels, those since the pandemic and at the moment amongst other items. (2)Participants send a hair strain to extract their hair cortisol and cortisone levels to compute their ratio. Data of N=553 (36% CN, 78% females, Mage= 26, SD=6.8) were analyzed. A food choice task was used in which participants rated healthiness and tastiness of different food items. Thereby, food choice sets were built for each participant with half of the trials showing foods where the tastier item was less healthy and thus self-control was required to make healthy choices. Results: To assess whether the COVID-19 pandemic led to increased stress, we compared the self-reported chronic stress-levels with prior literature using the same scale (TICS). We found that the mean reported TICS in our study was significantly higher than the score in the previous literature (MStudy=16.55, SDStudy=5.96 vs. Mliterature=9.4, SDliterature=6.3, $t(3019)=24.1$, $p<.001$). Further, we tested whether the ability to make healthier choices when self-control is required is linked to stress. We find a very low, albeit significant, correlation with pandemic-induced stress ($r(539) =-.09$, $p =.04$). When looking at the concentration of cortisol in relation to cortisone in hair, we also find a small negative marginally significant correlation with the percentage of healthy choices ($r(469)= -.09$, $p= .055$). Conclusion: We found that the higher chronic COVID-19-induced perceived stress, the lower people's ability to make healthy food choices when self-control is required. We found the same trend when using a biological marker of chronic stress, i.e., hair cortisol concentration. However, these negative correlations were very small and did not hold when investigating both countries separately. Together, our results showed much smaller effect sizes as compared to controlled lab studies investigating the impact of acute stress on self-control.

3-E-77 Cognitive reward control recruits medial and lateral frontal cortices, which are also involved in cognitive emotion regulation A coordinate-based meta-analysis of fMRI studies

Zarah Le Houcq Corbi[1], Felix Brandl, Satja Mulej Bratec, Christian Sorg
[1]Ludwig Maximilian University of Munich

Cognitive reward control (CRC) refers to the cognitive control of one's craving for hedonic stimuli, like food, sex, or drugs. Numerous functional magnetic resonance imaging (fMRI) studies have investigated

neural sources of CRC. However, a consistent pattern of brain activation across stimulus types has not been identified so far. We addressed this question using a coordinate-based meta-analysis of task-fMRI studies during CRC (22 studies with 741 subjects). To further characterize such a potential common CRC activation pattern, we extended our approach to three additional questions: (i) Do CRC meta-analytic results overlap with those during the control of emotional states, such as in cognitive regulation of aversive emotions (CER)? (ii) How does the control of motivational/emotional states link to the control of action states with less motivational/emotional valence such as in response inhibition paradigms, i.e., do meta-analytic result maps overlap? (iii) Does the control of motivational/emotional states constitute a consistent pattern of organized (i.e., coherent) ongoing or intrinsic brain activity? These questions were tested by conducting a coordinate-based meta-analysis of task-fMRI studies for CER (47 studies with 1445 subjects) and response inhibition (50 studies with 1042 subjects). Next, a contrast and conjunction analysis of the metanalytical results (CRC and CER as well as CRC/CER common activation and response inhibition) was performed. Finally, we conducted a seed-based whole-brain intrinsic functional connectivity (iFC) analysis in an independent data set of resting-state fMRI using the common CRC/CER pattern as seed. We found consistent CRC activation mainly in supplementary motor, dorsolateral prefrontal, and ventrolateral prefrontal cortices across studies. This activation pattern overlapped largely with CER-related activation, except for left-sided lateral temporal and parietal cortex activation, which was more pronounced during CER. It overlapped partly with activation during response inhibition in (pre-)supplementary motor, insular, and parietal cortices, but differed from it in dorsolateral and ventrolateral prefrontal cortices. Furthermore, it remarkably defined an iFC network covering activation patterns of both CRC and CER. Results demonstrate a consistent activation pattern of CRC across stimulus types, which overlaps largely with those of CER but only partly with those of response inhibition and constitutes an intrinsic co-activity network. These data suggest a common mechanism for the cognitive control of both motivational and emotional stimuli.

3-E-78 Neural oscillations implementing self-control in intertemporal choice

Georgia Eleni Kapetaniou[1], Marius Moisa[2], Christian Ruff[2], Philippe Tobler[2], Alexander Soutschek[1]
[1]Ludwig Maximilian University Munich, [2]University of Zurich

The dorsolateral prefrontal cortex (dlPFC) has been identified as one of the key brain areas implementing self-control in intertemporal choice. Little is known, however, about the neural oscillations in the dlPFC that are causally implementing self-control. With the present study we tested the effects of three different oscillation frequencies on intertemporal choice using noninvasive transcranial alternating current stimulation (tACS). Thirty-four healthy participants underwent tACS stimulation over the left dlPFC while performing an intertemporal choice task in which they had to make decisions between smaller-sooner and larger-later monetary rewards. In a within-subjects design, participants performed the task while undergoing alpha (10Hz), beta (20Hz), gamma (30Hz), and sham stimulation. While beta stimulation had no effects on decision making behavior, both alpha and gamma frequencies reduced impulsivity compared to baseline, such that participants were more likely to select the larger-later reward option under alpha and gamma stimulation. These effects were strongest around the individual indifference values. Here, we show that both alpha and gamma oscillations in the dlPFC implement self-control in intertemporal choice.

3-F-79 Perceptions of others' social affect and social cognition influences prosocial behavior

Leticia Micheli[1], Christina Breil[1], Anne Böckler[1]
[1]Leibniz University Hannover

OBJECTIVES: Gaze behavior plays an important role during social interactions. Yet, its immediate effect on interactive prosocial behavior remains unclear. We investigated whether perception of others' social understanding (via gaze-based cues during social interactions) affects the willingness to act prosocially. METHODS: 162 participants (Mage = 27.8 years, SD = 8.4, 41% female) witnessed six short video-based social interactions between two people, a Speaker and a Listener. Participants could hear (but not see) the Speaker telling either a negative or a neutral autobiographical episode (valence manipulation) and see the Listener, who looked either directly into the camera (direct gaze condition), downwards (averted gaze condition) or switched between direct and averted gaze (mixed gaze condition). After each social interaction, participants rated their perceptions regarding how empathic the Listener was (Empathy measure) and how much the Listener put themselves in the shoes of the Speaker (ToM measure). Subsequently, participants played a Trust Game, in which they were endowed with 10 chips and asked how many of these chips they would like to invest in the Listener. RESULTS: Results showed that the Listeners' gaze affected participants' trust in the Listener ($F(2,322) = 5.9$, $p = .003$). Listeners displaying mixed gaze were more trusted than listeners displaying averted gaze ($t(161) = 3.8$, $p < .001$). In addition, participants trusted more Listeners who heard a negative autobiographical episode ($F(1,161) = 26.7$, $p < .001$). Interestingly, these results were mediated by participants' perceptions of the Listener's social cognition skills. Listeners with mixed gaze, as well as Listeners who were told negative stories were considered more empathic and more understanding of the Speaker's mental states, and subsequently received greater investments in the Trust Game (indirect effect of gaze on trust through empathy and ToM: $b = -0.2$, $SE = .05$, $p < .001$; indirect effect of valence on trust through empathy and ToM: $b = .11$, $SE = .04$, $p = .012$). CONCLUSIONS: Results suggest that our perceptions of others' social cognition affect how much trust we place on them. To test the generalizability of these findings, a second experiment (in progress) investigates whether perceptions of others' social cognition also influences prosocial behavior in a non-strategic setting (e.g., giving behavior in a Dictator Game). Results of this second experiment will provide further information on whether the behavioral effects of perceptions of others' social cognition generalize across different prosocial domains or whether they depend on the existence of future interactions.

3-G-80 Human brain anatomy reflects separable genetic and environmental components of socioeconomic status

Hyeokmoon Kweon[1], Gökhan Aydogan, Alain Dagher, Danilo Bzdok, Christian Ruff, Gideon Nave, Martha Farah, Philipp Koellinger
[1]VU Amsterdam

Objective: The brain is the major locus of integration for genetic and environmental influences on behavior, including those related to socioeconomic status (SES). While neuroscience research has revealed a positive relation of SES with overall brain volume, as well as with regional cortical and subcortical volumes and cortical surface areas, little is known about whether such relationships are due to environmental or genetic effects. The primary goal of this study was to establish how SES relates to regional grey matter volumes (GMV) and differentiate known genetic from likely environmental components of these SES-GMV relations. Methods: We conducted a voxel-based morphometry analysis of GMV, using T1 MRI scan data from 23,931 adults from the UK Biobank (mean age 62, 57% female)

and composite measures of SES. We examined to what extent the genetic component accounts for the observed SES-GMV associations by comparing regression results of GMV on SES with and without controlling for a polygenic index (PGI) of SES constructed from multiple genome-wide association results (effective N ~ 900K). Furthermore, we studied the contribution of one lifestyle factor, body mass index (BMI), which is related to SES differences in nutrition, physical activity and obesity, to explaining the environmental component of SES-GMV relation. Results: We found widespread small, positive associations (partial R-squared not exceeding 0.5%) between SES and localized GMV. Controlling for the genetic underpinnings of SES reduced these SES-GMV associations on average by 38%. The relative reduction was largest bilaterally in the anterior insula, ventrolateral prefrontal cortex, and ventral striatum, suggesting that SES differences in these regions largely reflect a genetic basis. SES-GMV associations were less affected by PGI in the cerebellar, lateral temporal, and lateral parietal regions, implying that these regions are more prone to influences beyond genetic dispositions for SES. One such influence, correlated with SES, is BMI. On average, BMI accounted for almost a half of the residual SES-GMV associations after controlling for PGI, suggesting that lifestyle-related effects are likely to play an important mediating role for the effects of the socioeconomic environment on the brain. Conclusions: We disentangled possible genetic and environmental pathways by which SES affects brain structure. In particular, we found that some regions of the brain are linked to SES primarily due to underlying common genetic factors, while others are likely more susceptible to environmental influences.

3-G-81 Aging and reinforcement learning

Wei-Hsiang Lin[1], Marina Kunchulia[2], Michael Herzog[1]
[1]EPFL, [2]Free University of Tbilisi

Learning to achieve one's goal in a complex environment is a complicated task. In reinforcement learning (RL) tasks, an agent interacts with the environment to learn optimal actions. In humans, striatal areas are strongly involved in these tasks. During aging, neurotransmitter levels in these regions decrease provoking the question of how RL changes. Here, we developed a RL paradigm comprised of several different states. For each state, there were four actions, which brought the participants to another state. The objective of the participants was to reach a pre-defined goal state as many times as possible in 8 (short inter-stimulus interval condition) and 40 (long inter-stimulus interval condition) minutes. We tested a cohort of 40 healthy older participants (mean age = 68.75 ± 8.24) and 30 healthy young adults (mean age = 25.03 ± 4.09). Young participants reached more goals and made more correct actions, within the given time limit, than the ageing cohort. In addition, older participants exhibited more perseverative behaviors, namely, they repeatedly performed specific state-action pairs regardless of the correctness. By applying a Q-learning model, we observed a significant difference between older and young people in the exploration rate but not in the learning rate and forgetting rate. From the computation model, we demonstrated that the deteriorated performance for ageing is due to the deficit in neither learning nor working memory but from more suboptimal judgments of the action selections. These results are in line with the studies that less responsiveness of the striatal areas leads to more perseverative actions.

3-G-82 Age is associated with increased but more biased prosocial behaviour across the globe

Jo Cutler[1], Jonas Nitschke[2], Claus Lamm[2], Patricia Lockwood[3]
[1]The University of Birmingham, [2]University of Vienna, [3]University of Birmingham

Population ageing is a global phenomenon with significant implications across society. Ageing is associated with a multitude of neurobiological changes that have implications for social and economic decision making. Prosocial behaviours - actions that benefit others - promote mental and physical health across the lifespan and can save lives during the COVID-19 pandemic. METHODS. We examined whether age predicts prosociality in a preregistered global study (46,576 people aged 18-99 across 67 countries). Two acutely relevant measures provided a financial and a non-financial index of prosocial behaviour: i) distancing during COVID-19 and ii) a dictator game decision whether to donate to hypothetical charities. Participants stated the percentage of a specified amount of hypothetical money they would a) keep, b) donate to a national charity, and c) donate to an international charity, with both charities providing medical support for COVID-19. This allowed us to match the cause between the national and international charity and measure donation decisions independent of established habits. RESULTS. Strikingly, age positively predicted prosociality on both measures, with increased distancing and donations amongst older adults in the majority of countries across the globe. However, older adults were more biased than younger adults in choosing who to help, making larger donations to national over international charities and reporting increased ingroup preferences. Ingroup preferences helped explain greater national over international donations. Results were robust to internal replication and several control analyses, including accounting for subjective and objective personal wealth, country-level wealth and COVID-19 severity, perceived risk, and self-reported health. CONCLUSIONS. Older adults were more prosocial in their distancing and donation behaviours but also more biased than younger adults, preferring to help others in the same country, compared to people abroad. These findings have vital implications for predicting the social and economic impacts of ageing populations, increasing compliance with public health measures, and encouraging charitable donations.

3-H-67 Episodic decision-making via a process of cascading episodic sampling (CASES)

Achiel Fenneman[1], Alan Sanfey[2]
[1]Radboud University Nijmegen / Rhine-Waal University of Applied Science, [2]Radboud University Nijmegen

Daily life often presents us with complex decisions, the properties of which typically prevents us from relying on previous statistically-learned outcome values. Instead, we have to rely on our episodic memories to determine the actions that are available to us - and which of these is most attractive choice. In order to model this memory-based decision-making, we need to answer two related but independent questions: which memories do we rely upon, and how do we use these memories to make a decision. Though previous models and theories of memory-based decision-making exist, none provide a psychologically plausible answer to both of these questions. Here we describe such a model based on recent findings in memory and decision-neuroscience. This model is built around a bottom-up cascading spread of memory activation: the activation of a memory results in a simultaneous feed-forward and feed-back process. The feed-forward process generates evidence to either positively or negatively bias the action previously associated with this memory. If this process does not result in sufficient evidence to warrant a decision, then a feed-back pattern-completion process reinstates the features that make up the activated memories. These reinstated features then activate an additional set of memories, trigging the next feed-forward and feed-back processes, until a decision has been reached. We formalize this process in a model of Cascading Episodic Sampling (CASES), and show that this non-random sampling process is robust to large number of stored memories, and provides a faster decision speed as the number of memories increases. This model parsimoniously captures, and expands upon, a range of key findings and previous models in the topic of memory-based decision-making.

3-I-70 Parallel representation of context and multiple context-dependent values in ventro-medial prefrontal cortex

Nir Moneta[1], Mona Garvert[2], Hauke Heekeren[3], Nicolas Schuck[1]
[1]Max Planck Institute for Human Development Berlin, [2]Max Planck Institute for Human Cognitive and Brain Sciences, [3]Freie Universität Berlin

Objective: The ventromedial prefrontal-cortex (vmPFC/mOFC) is known to signal values that guide decisions. But how preferable available options are depends on one's current task. Goal-directed behavior, which involves changing between different task-contexts, therefore requires knowing how valuable the same options will be in different contexts. Whether such a multiplicity of values exists and how it might interact with value-free information about task structure in the same area is unknown. We tested whether multiple task-dependent values influence behavior and asked if they are integrated into a single value representation or are co-represented in parallel within vmPFC signals. Methods: Thirty-five participants alternated between tasks in which stimulus color or motion predicted rewards, while undergoing fMRI. On each trial, participants were cued to focus only on one context and had to choose the cloud leading to the largest reward. This meant that the same set of visual features was associated with different values in different task states. Although in each trial task states were associated with different values, across trials they had the same values and were dissociated from perceptual input. Results: Behavioral analysis indicated that participants did not merely perform a value-based choice among currently relevant values. Rather, both reaction times and accuracy showed they retrieved the values of irrelevant features and computed the resulting counterfactual choice. Using multivariate fMRI classifiers we investigated the probability distributions over values, which shed light on the complexity of representations in vmPFC. We found that the vmPFC maintains representations of the value associated with the current context/task state and the hypothetical value that would be expected in the alternative task state. Current task context could also be decoded from the same region. Crucially, we show these representations interact: expected and hypothetical values compete, and stronger task state representations are associated with stronger representations of its value. Finally, our results suggest a link between neural representations of task states and their associated values to their influence on behavior. Conclusions: We shed new light on vmPFC's role in decision making and bridge between previous perspectives of its role: first, it is involved in mapping observations onto a mental map, i.e. inferring the current task state; second, it is involved in computing the value expectations that result from the inferred state and third it also computes the hypothetical value that would result if one would have been in a different task state.

3-I-83 Neurocomputational processes of inferring others' preferences for information

Valentin Guigon[1], Julien Benistant[1], Marie Claire Villeval[1], Jean-Claude Dreher[1]
[1]CNRS

Information has an intrinsic value that is related to its capacity to reduce uncertainties about environmental contingencies. Humans seek information for instrumental, hedonic or cognitive motives and choose to approach or avoid information depending on the uncertainties they wish to clarify. This matters in social settings, in which one is often required to sample information about others and to infer their preferences. Estimation of others' preferences rely on computing the value one attributes to an item and to simulate other agent's value. Prior personal beliefs and others' beliefs about an item, drawn from past actions, are weighted and updated. We investigated the neural processes of inferring others'

preferences for information by adapting a sender-receiver game for two experiments with settings in which information have a cognitive utility. Prior to experiments, participants rated various political organizations from which we computed a social distance to each organization. In a first behavioral experiment, we elicited participants' beliefs regarding true and fake news via probability elicitation. We then asked them to disclose their desirability (Willingness-to-Pay) for each stimulus. In a second experiment, we used fMRI and elicited participants' beliefs for the same stimuli. We then asked participants to match the desirability of a receiver who did the first experiment. This receiver was drawn from a pool of 20 receivers. In a first Control condition, participants were provided no information about the receivers. In a second Cue condition, beliefs were elicited in participants about the receivers' beliefs. Each cue consisted in the receivers' social distance to an organization that was congruent with the news stimulus. Preliminary results show that participants are not better than chance level at evaluating true and fake news and that they are not displaying extreme biases in their veracity evaluations. Despite no difference with receivers in stimuli evaluation, senders seem to hold internal models of receivers' desirability that reflect neither their own beliefs regarding the news nor the receivers' behaviors. We will test senders' decision models of inferring the receivers' desirability for the news by taking into account veracity evaluation of the news, estimated population desirability for the news and beliefs about the receivers' beliefs. We expect senders to weight each parameter to take a sending decision. We will run model-based fMRI analysis and a Bayesian model to investigate how these parameters are weighed by the brain and to reveal how prior personal beliefs and others' beliefs about news are weighted and updated.

3-I-84 How robust is rationality?

Felix Nitsch[1], Tobias Kalenscher[1]
[1]HHU Düsseldorf

Study Objective: Neoclassic economic choice theory assumes that decision-makers make choices as if they were rational agents. This assumption has been critically challenged over the last decades, yet systematic aggregation of evidence beyond single experiments is still surprisingly sparse. Here, we asked how robust choice-consistency, as a proxy for rationality, is to endogenous and exogeneous factors, such as the influence of drugs, stress, aging, brain damage or other factors. Methods: To this end, we conducted a systematic quantitative literature research, reviewing 5327 articles, identifying 44 as relevant that contained hypothesis tests on possible influence factors of choice-consistency. To assess the evidential value of any effect of such influence factors on choice- consistency, we conducted a P-Curve analysis. P-Curve analysis combines the half and full p-curve method to make inferences about evidential value. In particular, if the half p-curve test is right-skewed with $p<.05$ or both the half and full test are right-skewed with $p<.1$, then p-curve analysis indicates the presence of evidential value. This combination test, introduced in Simonsohn, Simmons and Nelson (2015) 'Better P-Curves' paper, is much more robust to ambitious p-hacking than the simple full p-curve test is. Results: Descriptively, 78% of all significant p-values ($p<0.05$) fell in the lower half of the range of significant p-values ($p<0.025$). In line with this, our P-curve analysis indicates evidential value for that choice-consistency is affected by endogenous or exogeneous factors (binomial test: $p=.0053$; Stouffer method: $Z=-10.76$, $p<.0001$ for full p-curve and $Z=-11.4$, $p<.0001$ for half p-curve). Further, we find no evidence, that studies' evidential value is inadequate (binomial test: $p=.846$; Stouffer method: $Z=6.77$, $p>.9999$ for full p-curve and $Z=11.04$, $p>.9999$ for half p-curve). Additionally, the statistical power estimate of the included studies amounted to 96% (90%-CI: 91%-99%). However, due to the breadth of the contemporary research

agenda, the lack of replications and, often, the inaccessibility of original data in the field of choice-consistency, it is currently not possible to draw meaningful conclusions regarding specific influence factors beyond single experiments. Conclusions: Despite this lack of specificity, our results implicate that people's decisions might be a noisier and more biased indicator of their underlying preferences than previously thought. We conclude that rationality cannot be assumed unconditionally.

3-I-85 Is SAT Implemented by Adjusting the Frequency Threshold?

Marcin Penconek[1]
[1]University of Warsaw

Speed-Accuracy Trade-Off (SAT) has been widely studied; however, the underlying neuronal mechanism allowing the brain to implement SAT is not clear. Drift Diffusion Model (DDM) conceptualizes SAT through the selective influence hypothesis postulating that SAT effects are the result of adjusting decision threshold. Simulations of the Wang's model suggest that SAT is implemented as the dopamine dependent plasticity of cortico-striatal synapses (over longer periods of time) or through a constant top-down control signal (trial to trial adjustments) which affects the slope of ramping activity (Lo and Wang, 2009). SAT effects have recently been studied through a behavioral experiment (Rafiei and Rahnev, 2021) which suggests that DDM cannot account for the experimental effects assuming the selective influence hypothesis. This experiment shows a robust U-shape relationship between SAT levels and the difference in reaction time between error and correct trails. This kind of relationship is seen as a validity test for decision-making models. However, it can also be used to test alternative hypotheses on how the brain can implement SAT effects. Simulating the phenomenological network model with binary neurons (Penconek, 2020) reveals that the U-shape relationship can indeed be achieved by adjusting the frequency threshold in the winning attractor and no alternative hypothesis (such as control signal) leads to a similar result. This provides a computational argument in favor of the hypothesis that adjusting the frequency threshold is the main mechanism behind SAT. Frequency threshold can be implemented through adjusting inhibition level in the downstream circuits. Our analysis does not exclude the possibility that other neuronal mechanisms also interfere with SAT effects.

3-I-86 Preferential and perceptual context effects

Dimitris Katsimpokis[1], Laura Fontanesi[1], Jörg Rieskamp[1]
[1]University of Basel

Economic decisions are affected by their context: participants' revealed preferences for specific options can change when these options are embedded in different choice sets (e.g., Busemeyer et al., 2019; Huber et al., 1982; Simonson, 1989; Tversky, 1972). These so-called "context effects" defy traditional economic models (Luce, 1959). While participants have been extensively studied in the preferential domain, it has been recently shown that similar effects can be found in perceptual decisions as well (e.g., Trueblood et al., 2013). This discovery led to the development of computational cognitive models that assume common psychological mechanisms for context effects across perceptual and preferential choices (e.g., Trueblood et al., 2014). However, this assumption has not been tested within participants. We set up an experiment where participants make both preferential and perceptual decisions in a full-factorial within-subject design. We found no correlation between participants' choices across the preferential and perceptual domains. Finally, we evaluate whether a version of the race diffusion model (Tillman et al., 2020) captures the choice and response time patterns of participants.

3-I-87 A reinforcement learning race diffusion model accounts for effects of learning contexts on choices and response times

Laura Fontanesi[1], Amir Hosein Hadian Rasanan[2], Jamal Amani Rad[2], Sebastian Gluth[3], Jörg Rieskamp[1]
[1]University of Basel, [2]Shahid Beheshti University, [3]University of Hamburg

A recent wave of studies has proposed reinforcement learning models that can account not only for choice patterns but also for response times. In our previous work, we showed how different contexts in learning by feedback tasks affect not only choice proportion but also decision time: Participants tend to give faster responses in high-value contexts compared to low-value contexts. We explained such effects in two separate studies, in which we combined a traditional reinforcement learning model-which models across-trial dynamics-with a diffusion model-which models within-trial dynamics. However, it remains to be assessed whether the magnitude and sign of rewards are associated with different decision mechanisms. Moreover, the diffusion model cannot be generalized to decisions between more than two options. In this study, we collected data in a novel reinforcement learning task with decisions between two alternatives, in which we manipulated the magnitude and sign of rewards in a within-participant design. We found that the two manipulations had overall different effects on the joint choice proportion and response times patterns: the magnitude effect on response times is only present when the sign of the rewards is positive. We propose a new reinforcement learning race diffusion model in which a sigmoidal mapping of values to the accumulation rates takes into account both the context of a specific trial (by centering the sigmoid at the average values of the presented options) and the context of the experiment (by subtracting a reference point from the options values). We show how our model outperforms a previously proposed race model which uses a linear mapping and adds a constant to all accumulation rates (Miletić and colleagues, 2021). Our model, therefore, provides a concise and comprehensive account of value effects on decision-making in reinforcement learning while naturally generalizing to decisions between more than two options.

3-I-88 Economic valuation in the brain: An activation likelihood estimation meta-analysis of fMRI data involving a willingness-to-pay contrast

Alice Newton-Fenner[1], Danielle Hewitt[1], Jessica Henderson[1], Nick Fallon[1], Carl Roberts[1], Andrej Stancak[1]
[1]University of Liverpool

Background: Subjective valuation is a complex process incorporating a rich network of brain activations, and requires the amalgamation of an individual's perceptions, prior knowledge, and reward expectations of a given stimulus. It has often been implicitly defined through differing methodology; such as liking scales, unpleasantness ratings and binary 2AFC decisions. Subjective valuation can also be denoted as an economic value, measured using the metric of willingness-to-pay (WTP) in auction-based experiments. WTP has several advantages, including its ability to characterise value along a linear, discrete scale, its universal applicability to stimulus type (e.g. food/pain/music) as well as to both reward and punishment, and its ecological validity as a familiar, everyday metric. Theoretical understanding of subjective valuation would benefit from empirical investigation of brain activations for economic valuation, to address the confounding effect of task heterogeneity. The Becker-DeGroot-Marshack (BDM) auction was chosen for analysis as it is a demand-revealing mechanism widely used in value-based decision-making research as a tool to elicit WTP. Method: Twenty-two economic valuation studies utilising a BDM paradigm (703 participants; 164 foci) were selected by systematic review. Coordinate based activation likelihood estimation meta-analysis was performed correlating neurobiological

processing of auction items with the participants WTP. Additional contrast analyses were also performed comparing studies where the BDM was performed inside the scanner to those where the BDM was performed prior or post scanning, and comparing studies using food items as stimuli to those using non-food items. Results: WTP positively correlated with fMRI-BOLD activation in the bilateral ventral striatum. Contrast analysis revealed increased likelihood of activation for WTP for non-food items, relative to food items, in the left medial orbitofrontal cortex. No difference in brain activation patterns was found between outside vs inside scanner BDMs. Conclusion: Our findings offer succinct empirical support for the ventral striatum as the core structure responsible for the formation of economic value, irrespective of the type, quality or quantity of item. The results also suggest that the encoding of decision values for secondary-reward stimuli recruits greater cognitive resources in the left medial orbitofrontal cortex. When and where the BDM is completed during a study did not affect the WTP-related activation in bilateral ventral striatum, providing support for the robustness of BDM as a method for eliciting WTP.

3-I-89 People adaptively use information to alter internal and external states

Irene Cogliati Dezza[1], Christina Maher[1], Tali Sharot[1]
[1]University College London

Background Knowledge influences peoples' internal states and external reality. For example, information can impact internal states by inducing sadness, joy or fear and altering people's confidence in their ability to comprehend and predict the world around them. Information also alters external states by guiding people's decisions and actions. Thus, anticipating the impact of information on how we understand reality, feel about it and act upon it, is crucial for making information-seeking decisions that will increase survival and improve wellbeing. Here, we examine whether humans anticipate the impact of information on both their internal and external states and use those predictions to implement adaptive information-seeking behaviors. Methods A total of 210 participants in two experiments played a lottery game and on each trial could purchase more information about the lottery before deciding whether to gamble, or having a computer decide for them. We manipulate the impact of knowledge on internal states (mood and confidence) and external state (points earned) by altering the statistical properties of the lottery, such that the two were orthogonalized. Participants reported their subjective evaluations over their internal state (expected mood and uncertainty) and factors that influence their external states (their likelihood of making a lottery choice themselves). We then built models to predict participants' information purchasing decisions from their subjective estimates and/or statistical properties of the lottery. Results. Participants accurately anticipated how information will impact their internal and external state and used those predictions to guide information purchase. In particular, participants' information purchasing decisions were best predicted from a model that included participants' anticipation of how information will impact their mood ($\beta = 0.405 \pm 0.081$ (SE), $t = 4.99$, $p < 10^{-3}$) their uncertainty ($\beta = 0.431 \pm 0.053$ (SE), $t = 8.08$, $p < 10^{-3}$) and the instrumental utility of information ($\beta = 0.324 \pm 0.032$ (SE), $t = 10.23$, $p < 10^{-3}$). These considerations led participants to be happier ($t(53)=2.43$, $p < 0.05$), less uncertain ($t(53)=17.8$, $p < 10^{-15}$) and earn more points ($t(53)=3.96$, $p = 10^{-3}$) after purchasing information relative to when they did not. These findings were replicated over the two studies. Conclusion Our results show that people prospectively consider the internal and external impact of information, and integrate these estimates into adaptive information-seeking choices that improve their mood, confidence and external resources.

3-I-91 Evidence Accumulates for Individual Attributes during Value-Based Decisions

Douglas Lee[1], Todd Hare[2]
[1]Institute of Cognitive Sciences and Technologies, National Research Council of Italy, [2]University of Zurich

Traditional models of value-based decision-making are based on the principle that subjective values of individual choice options are calculated in the brain before being compared across options. In line with this, a plethora of studies have examined choice phenomena (e.g., choice accuracy/consistency and response time) through the framework of diffusion to bound models, most notably, the drift-diffusion model (DDM). We considered the prominent alternative hypothesis-- that the individual attributes of choice options are compared rather than their integrated overall values, and that the dynamics of choice deliberation are thus driven by individual attributes. Using simple variants of DDM (based on overall value, based on individual attribute measures, and a hybrid of the two), we demonstrate that individual attributes are more informative than overall value. However, we show that the hybrid model wins a formal model comparison, suggesting that overall value ratings provide useful information beyond that provided by the explicitly measured individual attributes. The hybrid model dominates across 6 previously published datasets. In particular, this model accounts for the published effects of attribute disparity (a formal measure of how dissimilar the options are in terms of attribute composition) on choice consistency and response time (in addition to the traditional effects of overall value difference).

3-I-93 Late integration of prior belief in drift-diffusion models

Stefan Bucher[1], Paul Glimcher[2]
[1]University of Tübingen / Max Planck Institute, [2]New York University

OBJECTIVE: Drift-diffusion models have been tremendously successful descriptions of the choice process. The abundant behavioral evidence they capture, however, mainly consists of choices and response times, which provide only limited identification of the underlying choice process. In order to examine when and how a prior modulates the choice process during evidence accumulation tasks, we designed a psychophysical experiment that can reveal the evolving decision variable. METHODS: Using a previously studied static perceptual choice task, we systematically varied the stimulus duration in an unpredictable manner, as well as the prior probability that either of the two choice alternatives was the correct one. In addition to subjects' binary choices, we also elicited their decision confidence in an incentive compatible manner. RESULTS: We found that for short trials the evolution of the log-odds of choice is precisely affine linear in log-time: Their dependence on the prior gradually decreases towards zero as time-in-trial evolves, as predicted by the biased starting point in a Bayesian model with exogenous stopping (i.e. in absence of stopping boundaries). However, we found that given sufficient time, conditional choice probabilities became independent of the prior, which is inconsistent with a biased starting point in a drift-diffusion model with endogenous stopping. While subjects' beliefs correlated strongly with their choices, we found that their beliefs did not incorporate the prior even when choices did. CONCLUSION: The two seemingly contradictory findings can be reconciled by a model in which the evolving decision variable captures only the likelihood of the accumulated evidence, rather than being modulated by the prior like a posterior belief. This accumulated evidence then appears to be combined with the prior only after evidence accumulation has stopped. This interpretation, which implies that subjects accumulate the same amount of information regardless of the prior, is further supported by our observation that decision confidence does not take the prior into account even when

choices do, suggesting that the prior is maintained by a cognitive process separate from evidence accumulation that subjects use only when evidence is scarce.

3-I-95 Brain functional connectivity predicts choice inconsistency

Asaf Madar[1], Vered Kurtz-David[2], Adam Hakim[1], Ido Tavor[1], Dino Levy[1]
[1]Tel Aviv University, [2]New York University

Objective: Humans are often inconsistent in their choices when choosing between bundles of goods. In an effort to explain this phenomenon, previous studies examined the neural correlates of inconsistent choices. Most studies find several value related brain areas to be the main correlates of inconsistency, but recent work in our lab suggests that motor brain areas also correlate with inconsistent choices. Here, we propose that choice inconsistency might also arise from alterations in the inter-region brain connectivity and not solely from activity within isolated areas. We use functional connectivity measures from resting state fMRI to predict subjects' inconsistency measures. We hypothesized that functional connectivity between value and motor areas will best predict choice inconsistency, and, therefore, compared performance by models solely based on these regions with other models based on whole brain parcellation. Methods: Thirty-nine subjects completed a task-free resting-state fMRI scan before performing a value-based risky choice task. For each subject, we computed several aggregate inconsistency measures based on her choices in the risk task. Furthermore, we computed two functional connectivity matrices for each subject: (1) hypothesis-driven, based only on value and motor regions of interest, (2) exploratory, based on a whole brain 100 parcel parcellation. As the whole brain connectivity analysis includes a very large number of features, we performed a feature selection step to include only the 45 features that were most correlated with the inconsistency measures. For the value and motor connectivity analysis, we used all 28 features. We then built two linear models with LASSO regularization to predict subjects' inconsistency measures from the selected features. Models were evaluated using leave-one-out cross validation. Results: We find a significant correlation between subjects' inconsistency measures and the predicted scores by the whole brain functional connectivity model ($r=0.356$, $p=0.042$). In contrast, the value and motor functional connectivity model did not produce a significant correlation ($r=0.279$, $p=0.116$). However, based on the coefficients of the whole brain model, we found that the most important features include the functional connectivity between motor, value, and attention related areas. Conclusions: These results further strengthen the notion that neural activity in both value- and motor areas contribute to choice inconsistency. Moreover, our results suggest that inconsistent choice behavior arises from general neural traits such as functional brain connectivity at rest, and not only from brain activity during actual choices.

3-J-71 Browsing under threat: high-level features of web searches altered during the pandemic and predicted population stress levels

Christopher Kelly[1], Bastien Blain[1], Tali Sharot[1]
[1]University College London

Objective: The 2020 pandemic generated a set of practical & mental challenges. To overcome these, people turned to the internet. Here, we examined how the high-level features of information people sought online (e.g. instrumental utility of information, valence) changed during the pandemic. To do so, we analyzed over a trillion internet searches submitted weekly before & during the pandemic in the UK & US. We related these measures to (i) stress levels reported weekly by 20K individuals & (ii) population-

level mobility data extracted from mobile phones. We theorized that during the pandemic people will be more likely to seek answers to questions that can guide action & increase comprehension of the world around them. Methods: We calculated the % of "How to" & "How do" Google searches submitted in the UK & US every week from 2017 to 2021. These searches likely result in information that can guide action (e.g. "how do you install Zoom?"). We did the same for "What" & "Why" searches. These likely increase comprehension (e.g. "what is Zoom?"). We also calculated the %positive minus %negative words entered for the most popular searches by matching words to an emotion lexicon. We then examined (i) how these features changed during the pandemic, & (ii) how these changes related to population stress levels reported weekly by 20K people in the UK. We dissociated the effects of stress from the effect of being confined to one's residence by controlling for mobility rates quantified from mobile phone data. Results: There was a sharp increase in proportion of searches that could guide action (i.e. "How to/How do") during the pandemic relative to before ($p = 0.0001$), which was positively associated with weekly population-stress levels, controlling for mobility rates ($p = 0.0001$). This suggests that under stress people seek information with high instrumental utility. This feature of web searches predicted population stress levels better than searches for specific stress-related terms (e.g. "anxiety"; $p = 0.003$). Mobility rates, but not stress levels, were associated with the proportion of "What/Why" searches, suggesting that home confinement may have increased the desire for general knowledge. Searches were also more negative. Results were replicated in UK & US. Conclusion: The increase in proportion of searches that can guide action may have contributed to human resilience in response to the pandemic. Markedly, this feature of web searches predicted population stress levels better than searches for specific terms (e.g. "psychiatrist"). An intriguing possibility is that the described method could be used for monitoring population stress levels beyond the pandemic.

3-J-97 Curiosity for information predicts well-being during COVID-19 Pandemic: contributions of loneliness and daily lifestyle

Annabel Losecaat Vermeer[1], Anne-Katrin Muth[1], Damiano Terenzi[1], Soyoung Park[1]
[1]German Institute of Human Nutrition, Potsdam-Rehbrücke

Objective: The COVID-19 pandemic confronted humans with high uncertainty, as enforced lockdowns severely disrupted people's daily social and health lifestyles, resulting in enhanced loneliness and lower well-being. Curiosity is central to human behavior, fostering adaptation in changing environments, besides impacting well-being. During the uncertain times of the pandemic, including plethora of pandemic-related information, curiosity and information-seeking as behavioral expression may particularly be critical to regulate mood and maintain well-being. Here, we first investigated which motives drive information-seeking. Second, we tested whether and how curiosity and information-seeking are related to well-being and mood (excitement, anxiety), and third, whether daily diet contributed to this relationship during the lockdown. Methods: To test this, 183 participants (M = 28.51 years) took part in a preregistered online study and completed questionnaires measuring trait curiosity, social and mental health, and rated their motivation to actively seek information that was, 1) potentially positive or negative (valence), and 2) concerned themselves or others (information target). Using a smartphone app, participants submitted their daily lifestyle ratings and daily momentary food intake for a week during lockdown. We extracted nutritive compositions from the submitted food intake, and performed linear (mixed) models and mediation analyses. Results: We found that information-seeking was driven by a significant valence by information target interaction ($p=.036$). Specifically, participants showed highest motivation to seek positive rather than negative information ($p<.001$), in particular

concerning themselves rather than about others (p<.001). Both trait curiosity and information-seeking positively predicted well-being, through reducing feelings of loneliness (indirect path: B=1.92, p<.001; B=0.08, p=.009, respectively). Moreover, individual's trait curiosity predicted averaged daily well-being (p<.001) and excitement (p<.05) a week into lockdown, but not anxiety (p>.05). Concerning diet lifestyle, only food intake with high-sugar predicted mood (anxiety, p<.01), but only in participants with relatively low, but not high, trait curiosity. Furthermore, we found that people with lower curiosity consumed food high in dopamine precursor tyrosine (p<.05). Conclusions: Curiosity may benefit well-being and mood in high uncertainty, by interacting with lifestyle measures (loneliness and nutrition), indicating the potential to modulate each of these players to increase well-being.

3-J-98 Effects of apathy and depression on mood and vigor

Huw Jarvis[1], Akshay Nair[1], Robb Rutledge[1]
[1]Yale University

-Objective Reward is known to influence both how fast we act and how we feel. These effects might be underpinned by distinct reward processes in the brain, governing the vigor with which we pursue rewards and the pleasure we derive from them. Clinical syndromes like apathy and depression are believed to affect reward processing, but it is unclear whether these symptoms relate to both mood and vigor. We investigated this question by testing a large smartphone-based sample (n=6,539). -Methods Participants used their personal device to play a fishing game, which required them to tap the screen successively to acquire rewards ('fish'). They periodically rated their subjective happiness on a slider, and completed questionnaires that measured their levels of apathy (Apathy Motivation Index) and depression (Patient Health Questionnaire). We examined mean latencies (i.e., tapping speed) and happiness ratings, and tested how these changed as a function of reward rate, which we systematically manipulated. -Results Apathy and depression both significantly predicted lower happiness during the task (Spearman's correlations: Apathy: rho = -0.21, p < 0.001; Depression: rho = -0.26, p < 0.001). Participants also reported greater happiness, and tapped faster, at high compared to low reward rates (Wilcoxon sign rank tests: Change in happiness: W = 1.6x10^6, p < 0.001; Change in latency: W = 1.2x10^7, p < 0.001). Replicating our previous results with lab and online samples, we found that the latencies of participants with greater apathy were more sensitive to fluctuations in reward rate (rho = -0.05, p = 0.02), but only in younger participants (younger participants (aged 40 and under): rho = -0.1, p = 0.005; older participants (aged over 40): rho = 0.01, p = 0.86). The extent to which changes in reward rate influenced happiness also decreased as a function of age (rho = -0.09; p < 0.001), and was also related to both apathy (rho = 0.06, p = 0.006) and depression (rho = 0.08, p < 0.001). -Conclusion This work highlights the potential for using smartphones to investigate clinical symptoms related to mood and motivation in the general population, and lays the foundation for future work to apply this technique in a clinical setting.

3-L-92 Consistency within change: Evaluating the psychometric properties of the predictive-inference task

Alisa Loosen[1], Tricia Seow[1], Tobias Hauser[1]
[1]University College London

Background: Rapid recognition and adaptation to sudden change is a hallmark of human behaviour. Many computational, neuroimaging and even clinical investigations capturing this ability have relied on

the same behavioural paradigm in which participants are asked to adapt to changes in the task environment (change-points; Nassar et al., 2010). However, the psychometric quality of this predictive-inference task has not sufficiently been examined, leaving unanswered whether this task is indeed suited to capture behavioural differences on a within- and between-subject level. Methods: We conducted a large-scale (N=330), re-test online study in which participants played this predictive-inference task on two occasions (~3 months apart). We assessed the internal (internal consistency) and temporal (test-retest reliability) stability of the task's relevant raw measures (i.e., confidence and learning rate) and their associations with a Bayesian learner. We estimated the internal consistency of these measures using their Spearman-Brown corrected Pearson correlation (rSB) and their test-retest reliability with their intraclass correlation (ICC). All psychometric scores were computed relative to the change-points in the task to ascertain whether measures robustly captured responses to environmental changes. Results: We show that the main measures capturing flexible action and belief adaptation yield good internal consistency (rSB≥0.72) and mostly moderate to good test-retest reliability (ICC≥0.65). However, links between these measures as well their association with the Bayesian learner lack convincing psychometric quality. Conclusions: We found that the main measures of this predictive-inference task consistently and reliably capture belief and behavioural adaptations before and after environmental changes, making them suitable for studies investigating individual differences. We also note that the complex links between these task measures and Bayesian predictions are of mostly low psychometric quality and should only be used with caution. Our findings have implications for the large corpus of previous studies using this task and provide clear guidance as to which measures should and should not be used in future studies.

3-M-90 A neuro-computational study of moral contagion

Julien Benistant[1], Valentin Guigon[1], Alain Nicolas[2], Jean Claude Dreher[1]
[1]CNRS, [2]Centre Hospitalier le Vinatier

People moral compass is often influenced by others' behavior. For example, a salesman whose colleagues are all dishonest to customers can end up mimic their behavior. In a new social environment, people may be more uncertain about the moral conduct of others. One way to learn the local norm is simply by observing how others behave. This implies that moral contagion, the process by which one align his moral behavior to fit the others' one, is dynamic. Here, we investigate the neuro-computational correlates of moral contagion. By combining a model-based approach with fMRI, the goal of this project was to understand both the computations occurring dynamic moral contagion as well as the brain regions engaged with these computations. We designed a new fMRI experiment based on a repeated binary-reporting task in which participants (n= 32 (15 males), mean age= 22.9 y) had to report the outcome of a dice roll. After observing the outcome, participants had to report what they saw from two options. One option was the true outcome and the other was an incorrect outcome. Each choice was associated with a payoff and choosing the incorrect outcome was always more profitable than reporting truthfully. This task allowed us to assess participants' misreporting behavior and to investigate how it evolved with time. The exposition to others' behavior was done using prediction periods. In these periods, participants had to predict what was the report of an individual who did the task in a previous session. On each trial, this individual was randomly selected from a group of 10 past participants. A correct prediction was rewarded with ?2, while an incorrect one was not rewarded. We varied within-subject and between blocks, whether participants had to predict or not the behavior of others as well as the degree of (dis)honesty of the others. In blocks with prediction periods, participants did alternatively

one period of the reporting task and one period of the prediction task. Preliminary results show that participants are influenced by others' behavior but with important inter-individual differences. A model integrating both a utility model and a Bayesian preference learning model is the best suited to explain participants' decisions as well as learning over the course of the task. This allows us to directly study the dynamic aspect of moral contagion. Preliminary fMRI results suggest that moral contagion leads to a modulation of the vmPFC activity, correlating with the size of the earnings associated with misreporting. Together, these findings provide insights for the neuro-computational bases of moral contagion.