Society for NeuroEconomics
NEUROSCIENCE • PSYCHOLOGY • ECONOMICS

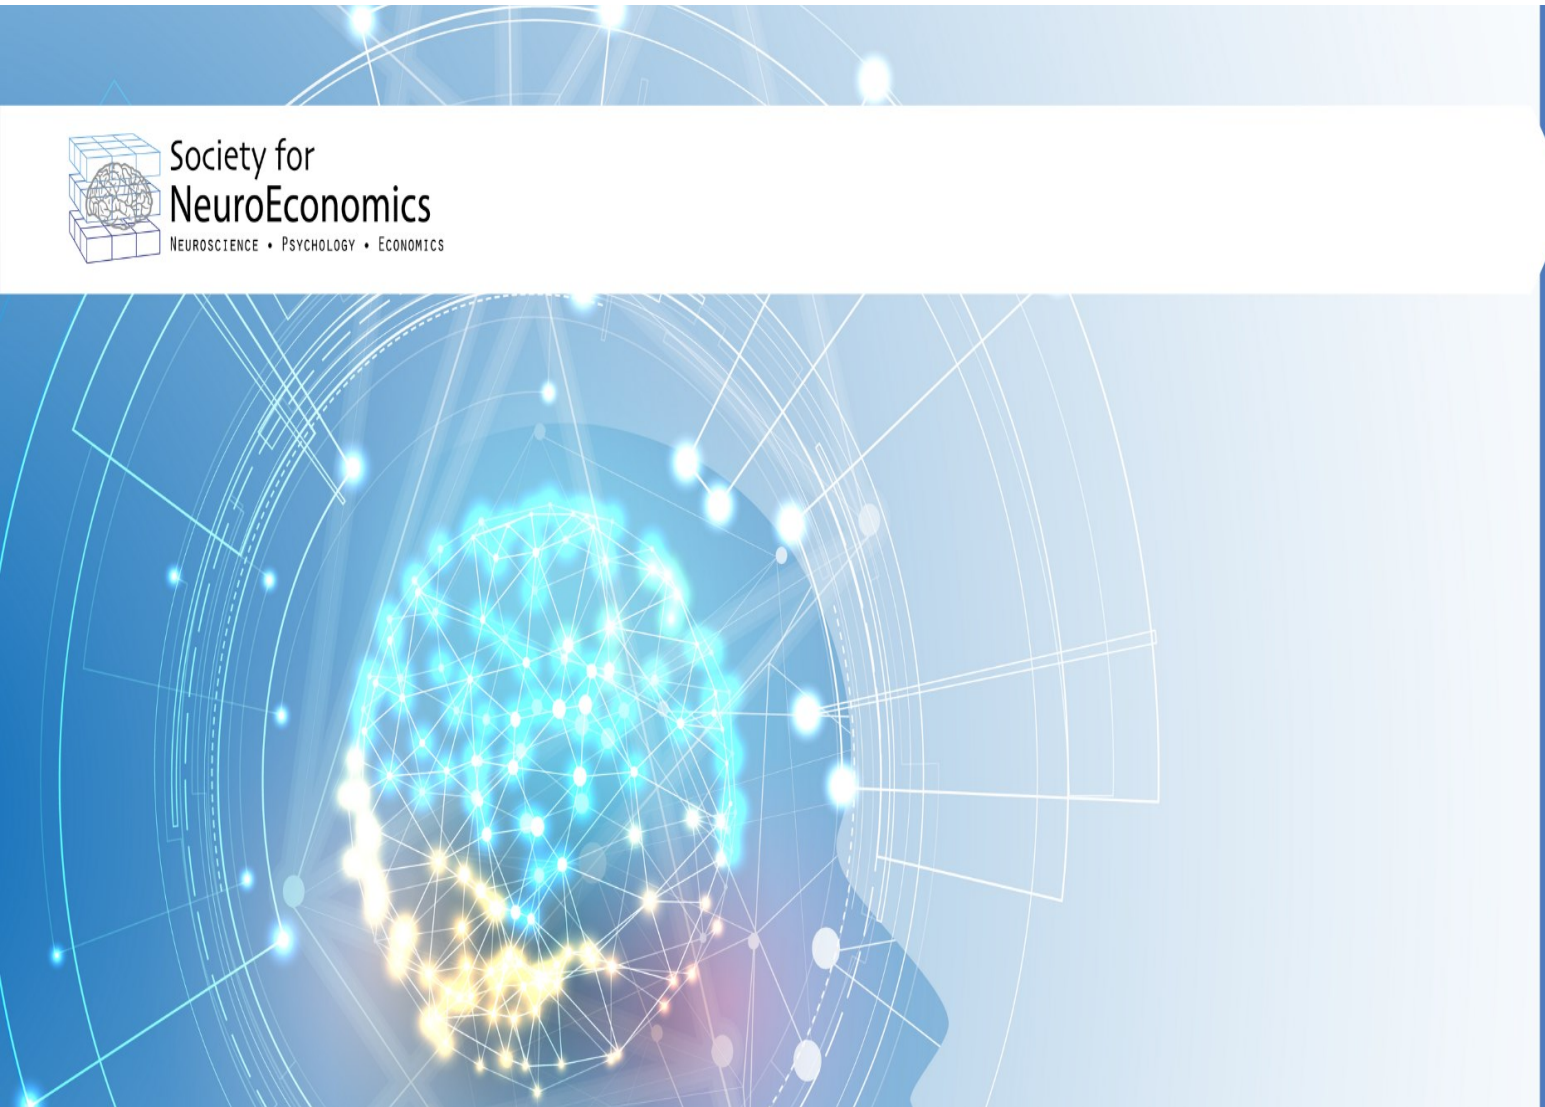

Abstract Book

vSNE

18th Annual Meeting

October 7 - 9, 2020

neuroeconomics.org | #VSNE2020 | @SocforNeuroecon

Contents

ORAL PRESENTATION ABSTRACTS

Session I: Risk, uncertainty, and delay

Considering what we know and what we don't know: Expectations and metacognition guide value integration during economic choice

Romy Froemer[1], Frederick Callaway[1], Amitai Shenhav[1]
[1]Brown University

Objective: Most models of economic choice assume that choices are at their core driven by noisy samples of value, but how people accumulate these samples remains unclear. While recent work has demonstrated that overt attention can shape the way in which option values are integrated into a decision, less is known about the extent to which this integration depends on prior expectations and beliefs about those values. Here we use a novel, sequential presentation paradigm to investigate how value integration is informed by what the decision-maker knows at a given point in time (predictive inference), and by how reliable they believe this information to be (metacognition). Method: Participants (N = 30) chose between pairs of options previously rated on their value, and the participant's confidence in those values. Options alternated on the screen, and we independently varied the order, value, and duration of item presentation. Participants could choose either option at any time prior to a 5 second deadline. We compare participant behavior to simulations of a Bayesian model that optimally integrates evidence from samples whose variance depends on both attention and confidence. Results: We found that participants' choices were shaped both by predictive inference and metacognition, and that this can be accounted for by a Bayesian model of value integration. Consistent with the integration of prior information, (1) choices display hallmarks of reference-dependent processing of the options' values and of the first seen item in particular, and (2) participants selected certain over uncertain options when overall value was high, but vice versa when overall value was low. Consistent with uncertainty-weighted updating, participants were faster and more accurate when their confidence in both options' values was higher. However, participants who were more confident on average showed faster but less consistent choices. We show that this maladaptive behavior is consistent with Bayesian integration using incorrect variance estimates (i.e. faulty metacognition). Conclusion: Our findings provide a more complete understanding of the information that shapes our evaluation of choice options. Our empirical data and modeling show that this information includes prior beliefs about value distributions, as well as beliefs about the precision of value representations. In so doing, our work demonstrates how growing research on Bayesian inference across domains of psychology and neuroscience can be leveraged to offer novel insights into economic choice.

Predicting Risk Attitudes from the Precision of Neural Magnitude Representations

Miguel Garcia[1], Gilles de Hollander[1], Marcus Grueschow[1], Rafael Polania[2], Michael Woodford[3], Christian Ruff[1]

[1]University of Zurich, [2]ETH Zurich, [3]Columbia University

Risk aversion occurs when a decision maker prefers a choice option with a smaller expected monetary payoff but more outcome certainty over an alternative option with a larger expected payoff but higher outcome uncertainty. This phenomenon is classically explained as reflecting properties of outcome valuation, i.e., the concavity of the utility function mapping monetary amounts to subjective utility. However, this explanation cannot account for the prominent empirical findings that subjects are risk averse even for very small gambles, and that they often tend to choose differently across repetitions of the same choice problem. To overcome these problems, it has recently been proposed that risk aversion in small-stake gambles may not reflect properties of valuation but rather of numerical cognition - how potential payoffs are mentally represented (Khaw et al., 2019). These proposals are based on the observations that humans (a) underestimate larger magnitudes, especially under uncertainy and (b) are stochastic in decision-making about numerical magnitudes. Since risky options are typically associated with larger payoffs, this underestimation bias may lead to risk aversion simply as a result of inaccurate and noisy perception. Crucially, such a perceptual account predicts that an individual's observed risk attitude should depend on the precision of the perceptual representations of payoffs. Here, we test this novel prediction both behaviorally and neurally. Subjects performed a perceptual numerical decision-making task during fMRI as well as a risky-choice task with different presentation formats for potential payoffs, i.e., non-symbolically as a pile of coins and symbolically as Arabic numerals. We fitted a computational model of noisy logarithmic coding (NLC) of numerical magnitudes (Khaw et al., 2019) that can jointly explain behavior in both task domains and quantified noise and bias parameters. Our results show that (1) the precision of mental representations of numerical magnitudes and payoffs is correlated within subjects, across perceptual and risky choices; (2) external noise due to presentation format leads to increased risk aversion; and (3) subjects with more reliable magnitude representations in parietal cortex during the perceptual task show less variable behavior in the risky-choice paradigm and are less risk-averse. In sum, our results show that risk preferences can be predicted from the precision with which numbers are (neurally) represented. More generally, our results highlight that aspects of economic behavior may be determined by capacity limitations in perceptual processing rather than properties of neural valuation.

From Value to Saliency: Neural Computations of Subjective Value under Uncertainty in Combat Veterans

Ruonan Jia[1], Lital Ruderman[1], Charles Gordon[1], Daniel Ehrlich[1], Mark Horvath[1], Serena Mirchandani[1], Clara DeFontes[1], Steven Southwick[1], John Krystal[1], Ilan Harpaz-Rotem[1], Ifat Levy[1]
[1]Yale University

Objective: Military personnel engaged in combat are vulnerable to Posttraumatic Stress Disorder (PTSD), following traumatic experiences in the battlefield. Prior research has mostly employed fear-related paradigms to unravel neural underpinnings of fear dysregulation related to PTSD. The ability to acquire and update fear responses depends critically on the individual?s ability to cope with uncertainty, yet the role of individual uncertainty attitudes in the development of trauma-related psychopathology has hardly been examined. Here, we investigated the association between PTSD-related alterations and the subjective valuation of uncertain outcomes during decision-making. Methods: We used a monetary gambling paradigm to explore the neural markers of both vulnerability and resilience to PTSD in combat veterans (24 with current PTSD, 34 controls). Participants chose between a certain gain (or loss), and playing a lottery which offered a larger gain (or loss) but also chance of zero outcome. Outcome probabilities for half of the lotteries were precisely known, and were ambiguous for the other half. fMRI was used to track neural activation while subjects completed 240 decisions. One choice was randomly picked for payment to ensure task engagement. We evaluated PTSD symptoms by CAPS (Clinician-Administered PTSD Scale). Results: Using a dimensional approach, we replicated our previous finding (Ruderman et al. 2016) that veterans with more severe PTSD symptoms were more averse to ambiguous losses (Pearson?s correlation $r=-0.30$, $p<0.05$). We additionally found that they were more averse to risky gains ($r=-0.39$, $p<0.01$). fMRI activity in the ventromedial prefrontal cortex (vmPFC) during valuation of uncertain options was associated with PTSD symptoms ($p<0.001$, corrected at FWE=0.05), an effect which was specifically driven by numbing symptoms. Moreover, the neural encoding of the subjective value of those uncertain options was markedly different in the brains of veterans diagnosed with PTSD, compared to veterans who experienced trauma but did not develop PTSD. Most notably, veterans with PTSD exhibited enhanced representations of the saliency of rewards and punishments in the neural valuation system, especially in ventral striatum, compared with trauma-exposed controls. Conclusion: Our results point to a link between the function of the valuation system under uncertainty and the development and maintenance of PTSD symptoms, and stress the significance of studying reward processes in PTSD.

Predicting individual differences in delay discounting based on fMRI

Leonie Koban[1], Arthur Lee[2], Daniela Schelski[3], Marie-Christine Simon[3], Caryn Lerman[4], Bernd Weber[3], Joseph Kable[2], Hilke Plassmann[1]
[1]INSEAD & ICM, [2]University of Pennsylvania, [3]Universitaet Bonn, [4]University of Southern California

Objective: Choosing between immediate but smaller rewards and future larger rewards is the basis of many everyday decisions. Delay discounting--the preference of immediate over future rewards--has been proposed as an individual trait predictive of important life outcomes and as a potential transdiagnostic marker of psychopathology and addictive disorders. Previous studies have investigated structural and resting-state correlates of delay discounting. Which functional processes during decision-making reliably predict interindividual differences in how much people discount future rewards remains unclear and is the central question of this research. Methods: We used machine-learning to predict individual differences in delay discounting based on functional brain activity patterns. In two independent fMRI studies (total N=257), adult volunteers performed an intertemporal choice task, making incentive-compatible choices between smaller immediate and larger later (LL) rewards. These choices allowed us to compute the hyperbolic discounting factor (k)--a standard measure of individual differences in delay discounting. We computed individual contrast images for choice onset, its parametric modulation by LL amount, and its parametric modulation by delay. The combination of these three contrast images from Study 1 (N=112 male participants, Bonn University, Germany) was used for 10-fold training (using LASSO-PCR) and cross-validation to predict individual differences in log(k). Study 2 (N=145 male and female participants, University of Pennsylvania) was used as an additional, independent test set. Results: We obtained a significant cross-validated prediction-outcome correlation in Study 1 ($r = 0.45$, $p < 0.001$). The resulting pattern also significantly predicted individual differences in discounting in Study 2 ($r = 0.47$, $p < 0.001$), demonstrating generalizability of the brain marker to a completely independent dataset, with different sample and task characteristics. Significant voxel weights ($q < 0.05$ FDR corrected) were found in ventromedial and dorsolateral prefrontal cortex, insula, and hippocampus, among other regions. Conclusions: Using machine-learning, we identified cross-validated and generalizable brain activity patterns that predict individual differences in delay discounting and could be employed as a brain-based measure of discounting in future studies. They point at the interplay of several brain systems driving individual differences in intertemporal decision-making, including the brain's valuation system, but also areas associated with conflict-processing, memory, and cognitive control.

Session II: Social decision making

The computational basis of individuals' learning under uncertainty in groups with collective goals

Sensen Song[1], Hanbo Xie[2], Marrten Speekenbrink[3], Jianlei Zhang[4], Xiaoxue Gao[2], Xiaolin Zhou[2]
[1]Tongji University, [2]Peking University, [3]University College London, [4]Nankai University

Despite the ample research on how individuals learn to exploit rewards and explore information under uncertainty, little is known about how individuals learn under uncertainty in

groups with collective goals. To address this issue, we developed a group-based learning task, in which participants performed Iowa Gambling Task (IGT) together as a group. We recruited groups consisting of five participants (118 groups, N=590) and scheduled them into a two (NoCommunication vs. Communication)-by-two (NoLeader vs. Leader) between-subject design. We additionally recruited 39 individuals who performed IGT alone and served as the baseline (Individual learning). On each trial, group members individually noted their preferred choice and then the group made a collective decision. Groups without leaders followed the majority rule to reach agreements, while groups with leaders made final decisions by the leaders after the presentation of group votes. Groups with communication were allowed to discuss for two minutes after every 20 trials. Results showed that the existence of both leader and oral communication enhanced group performance, but with no significant interaction between the two variables. Using computational modeling, we found, for Individual learning, that the Value plus Sequential Exploration (VSE) model, which integrates exploitation and exploration in learning, performed best to explain participants' choices among all existing models for IGT. Based on this model, we developed a social-influence-based VSE model for individuals' decisions in the groups, in which group members' choices would increase other members' choices of the same option in the next round of decision, as reflected by a parameter βother. Moreover, individuals in groups applied discount rewards (weighted by the same parameter βother) arising from collective choices to update the exploitation and exploration value of options when their own decisions are inconsistent with the group. Comparing parameters across different group conditions, we found that members in groups with communication received higher social influence from other members (βother) and updated exploration value to a larger extent compared to members in groups without communication. Furthermore, compared with members in NoLeader groups and non-leading members in Leader groups, leaders in Leader groups were more likely to be influenced during the task. Our model reveals the computational basis of individuals' learning under uncertainty in groups with collective goals, with social influence as a crucial factor in modulating individuals' exploitation and exploration strategies.

Interoceptive sensitivity is associated with neural representation of others' rewards

Luis Sebastian Contreras Huerta[1], Michel-Pierre Coll[2], Hongbo Yu[3], Annayah Prosser[4], Jennifer Murphy[5], Patricia Lockwood[6], Geoffrey Bird[1], Molly Crockett[7], Matthew Apps[6]
[1]University of Oxford, [2]McGill University, [3]University of California Santa Barbara, [4]University of Bath, [5]Royal Holloway, University of London, [6]University of Birmingham, [7]Yale University

Interoception -the ability to sense one's own bodily states- has been linked with multiple social processes, including empathy and prosocial behavior. Interoception is also crucial for reward evaluation, informing embodied experiences of reward anticipation and receipt that facilitates

goal-directed behaviour. Both interoception and social reward processing have been linked with activity in anterior insula (AI) and anterior cingulate cortex (ACC). Here, we examined how interoception relates to the neural processing of others' rewards in these regions, and whether they play different functional roles. METHODS: We measured in 59 subjects (i) interoception, using a task that quantifies people's sensitivity to internal bodily signals relative to external information; (ii) motivation to help others, using a task where people trade-off effort against profit for self and other; and (iii) neural reward representations, using an fMRI task where participants passively witnessed monetary gains and losses for self and other. Using representational similarity analysis on AI and ACC, we computed similarity between self/other reward representations in these regions. Thus, we tested the relationship between interoception and social reward processing along two different dimensions: (i) motivational, how incentivised people are when working to get others' rewards, and (ii) neural, how similar the representations of self/other rewards are in AI and ACC. RESULTS: We dissociated two distinct links between social reward processing and interoception. Higher interoceptive sensitivity was associated with (i) being more motivated by others' rewards when choosing whether an effort is worth it and (ii) more dissimilar neural reward representations for self and other in right AI but not in ACC. However, participants who were more motivated to exert effort for others' rewards, regardless their interoceptive sensitivity, showed more similar BOLD activation patterns between self/other rewards in ACC but not in AI. CONCLUSIONS: People with higher interoceptive sensitivity are more motivated to exert effort for others' rewards. Importantly, AI and ACC seem to play distinct roles in integrating this information. Dissimilarity in right AI representations for self/other rewards was associated with higher interoceptive sensitivity, consistent with its role in self-awareness and self/other differentiation, crucial for social cognition. On the other hand, more similar self/other reward representations in ACC was associated with more willingness to help others. Thus, while ACC representations of others' rewards are tied to prosocial behaviour and motivation, AI representations are embodied.

Modeling models of others' mental states: characterizing Theory of Mind processes during cooperative interaction

Tessa Rusch[1], Saurabh Steixner-Kumar[2], Prashant Doshi[3], Michael Spezio[4], Jan Gläscher[2]
[1]Caltech, [2]University Medical Center Hamburg-Eppendorf, [3]University of Georgia, [4]Scripps College

Humans are distinctly skilled at cooperation, i.e., at acting in accordance with the perceptions, goals, and beliefs of others to facilitate own and others' gains equitably. To successfully cooperate, humans have to combine predictions of their partners' behavior with their knowledge of the environment and act according to the combined requirements of the interactive situation. Cognitively, this requires Theory of Mind (ToM): the capacity to estimate and represent others' mental states and predict rational behavior based on these mental states.

Here, we investigated these processes during real-time reciprocal coordination between two players engaging in a cooperative decision game while Electroencephalography (EEG) was recorded from both players. The game took place in a noisy and unstable environment and information was distributed asymmetrically among the players. To succeed, players had to model the dynamically changing environment and their partner's belief about it and integrate both pieces of information into a coherent decision. Thereby, the game combines social and non-social learning into a single decision problem. To quantify the learning processes underlying participants' actions, we modeled behavior with Interactive Partially Observable Markov Decisions Processes (I-POMDP). The I-POMDP framework extends single agent action planning under uncertainty to the multi-agent domain by including intentional models of other agents. Using this framework, we successfully predicted interactive behavior. Furthermore, we extracted participants' beliefs about the environment and their beliefs about the mental states of their partners, giving us direct access to the cognitive operations underling cooperative behavior. By relating players' own beliefs with their partner's model of themselves, we show that dyads whose beliefs are more aligned coordinate more successfully. Further, using the extracted set of beliefs, we computed players' trial-by-trial expectations about the environmental dynamics and their representations of the partners' expectations. Model-based EEG analyses showed that frontal theta codes for violations of players' own expectations. Interestingly, similar neural signals also coded for the presumed violations of the partner's expectation in the absence of experienced surprise. This comprehensive quantification of ToM processes during cooperation suggests that behavioral coordination relies on mental alignment.

Emotion Prediction Errors Guide Socially Adaptive Behavior

Joseph Heffner[1], Jae-Young Son[1], Oriel FeldmanHall[1]
[1]Brown University

How do we learn to make adaptive decisions, such as whether to avoid a risky financial endeavor or start a collaboration with a new colleague? A rich literature on value-based decision-making illustrates that choices are made based on the expectation of rewards, and that violations of these reward expectations--i.e., prediction errors (PE)--enable an agent to update their knowledge about their environment to facilitate survival. In parallel with research linking reward to decision-making, a separate literature also demonstrates that emotion exerts a powerful influence on people's choices. Although there has been interest in understanding how anticipated emotions affect behavior, relatively little work has examined how the violation of expected emotions influences decision-making, especially in the context of complex social interactions. To test how strongly reward and emotion PEs impact social behaviors such as punishing or helping others, we quantify how emotional experiences and their violations bias choices during multiple interactive economic games. We therefore developed a technique that measures real-time fluctuations in emotions as the decision-process unfolds, enabling us to

precisely map the subjective experience of emotion alongside economic rewards during social exchanges. Across three separate experiments, participants (N=789) played one of two behavioral economic games, capturing punitive responses to fairness violations, while simultaneously rating their affective experiences using a novel measure on both arousal and valence dimensions. By measuring emotions as a social interaction unfolds, we can compute the difference between emotion expectations and compare them to the actual emotional experience, effectively capturing emotion PEs. Across three studies, we reveal that emotion and reward have independent and dissociable contributions to choice, such that emotion PEs exert the strongest impact on deciding to punish or help another. Specifically, we found that valence PEs have a significantly stronger impact on motivating punitive choices than reward PEs. Using the temporal trajectories of the emotional responses in study two, we additionally find that these decisions can be decoded from their emotional responses as early as 430ms, suggesting emotions swiftly influence choice. Finally, in study 3, we demonstrate a functional dissociation between emotion and reward PEs as individuals with depression exhibit selective impairments in using emotion--but not reward--PEs. By demonstrating the power of emotion PEs in guiding social behaviors, these findings challenge standard decision-making models that have focused solely on reward.

Session IV: Valuation and value systems

Grid-like codes for novel inferences during value-based decision making

Seongmin Park[1], Douglas Miller[1], Erie Boorman[1]
[1]University of California, Davis

Generalizing previous experiences to construct values for current decision making in novel situations is a hallmark of behavioral flexibility. It has been hypothesized such flexibility depends on a cognitive map - a representation of relationships between states or entities in an environment. In particular, a cognitive map could theoretically allow for novel inferences without direct experiences that guide decision making in new situations. However, how cognitive maps are used to guide such novel value-based decisions is unclear. Here, we developed a novel task to address this question. During behavioral training participants learned the relationship between individuals in two independent social hierarchy dimensions through a series of comparisons between pairs at neighboring ranks in one dimension at a time. Without seeing the overall graph structure, participants could infer the latent social hierarchy structure through transitive inferences. During fMRI participants made novel decisions that required them to combine two separately learned dimensions to compute decision values. In each trial, participants were asked to choose the better partner for a given 'entrepreneur' (A) between two potential business partners (B and C) by comparing the 'growth potential (GP)' of two partnerships (A-B and A-C). The GP, the decision-value in the current study, is not determined

by the individual ranks in the social hierarchy; instead, it depends on the area drawn by two individuals over the abstract 2-dimensional (2-D) space. First, we find that discretely sampled abstract relationships between people are reconstructed into a unitary 2-D cognitive map in the patterns of activity in the hippocampus (HC) and entorhinal cortex (EC). Second, we find that the human brain utilizes a hexadirectional grid-like code in the EC, medial prefrontal cortex (mPFC), posterior cingulate cortex (PCC), temporoparietal junction (TPJ) and superior temporal sulcus (STS) for inferred direct vectors between entities in the reconstructed 2-D space. This grid-like signal is consistent between sessions acquired more than a week apart. Moreover, the neural grid-like codes in EC explain neural decision value computations in TPJ and mPFC during decisions. Finally, during decisions the EC and vmPFC activity also encodes the relative decision value, suggests their interplay in comparing the values constructed from the multidimensional cognitive map. Our findings show that a grid-like code is extended to encode trajectories in general abstract task spaces to construct values for novel decision making and further suggest a general mechanism underpinning flexible decision making and generalization.

Values Encoded in Orbitofrontal Cortex Are Causally Related to Economic Choices

Weikang Shi[1], Sébastien Ballesta[1], Katherine Conen[1], Camillo Padoa-Schioppa[1]
[1]Washington University in St. Louis

In the 18th century, Daniel Bernoulli, Adam Smith and Jeremy Bentham proposed that economic choices rely on the computation and comparison of subjective values. This hypothesis continues to inform economic theory and experimental research, but behavioral measures are ultimately not sufficient to prove the proposal (Camerer et al. 2005). Consistent with the hypothesis, when agents make choices, neurons in the orbitofrontal cortex (OFC) encode the subjective value of offered and chosen goods. Moreover, neuronal activity in this area suggests the formation of a decision (Padoa-Schioppa and Conen 2017). However, it is unclear whether these neural processes are causally related to choices. More generally, the evidence linking choices to value signals in the brain remains correlational. Here we show that neuronal activity in OFC is causal to economic choices. Causal links between a neuronal population and a decision process are demonstrated if one can predictably bias choices using electrical stimulation. Building on this concept, classic work established the causal role of the middle temporal area in motion perception, by showing that low-current stimulation biases (Salzman et al. 1990) while high-current stimulation disrupts (Murasugi et al. 1993) perceptual decisions. One challenge in using this approach for economic decisions is the lack of columnar organization in OFC. Since neurons associated with different goods are physically intermixed, it is impossible to selectively activate neurons associated with one particular good using electrical stimulation. We developed two experimental paradigms to circumvent this challenge. In Exp.1, monkeys chose between two juices offered sequentially. High-current stimulation (≥100 μA)

delivered during offer1 or offer2 consistently biased choices against the juice offered during stimulation. Furthermore, high-current stimulation during offer2 (when values were compared) increased choice variability. In Exp.2, we took advantage of the fact that offer value cells in OFC undergo range adaptation. Thus, for any juice, the increase in offer value induced by low-current stimulation should be proportional to the value range. If two juices are offered, stimulation should increase both values, but ultimately bias choices in favor of the juice with the larger value range. In the experiments, monkeys chose between two juices offered simultaneously, and we delivered electric current (50 μA) in half of the trials. Confirming our predictions, stimulation induced a choice bias strongly correlated with the difference in value ranges. Taken together, these results demonstrate that values encoded in OFC are causal to economic choices.

Retrieval-Constrained Valuation: Toward Prediction of Open-Ended Decisions

Zhihao Zhang[1], Shichun Wang[1], Maxwell Good[1], Siyana Hristova[1], Andrew Kayser[2], Ming Hsu[1]
[1]University of California, Berkeley, [2]University of California, San Francisco

Real-world decisions, such as how to spend one's evening or choosing a career path, are often open-ended, with goals, choice options, or evaluation criteria conceived by decision-makers themselves. These requirements, however, are largely absent from classical models of decision-making, thereby limiting their predictive scope. Here we take a step toward addressing this issue by expanding current models of decision-making and developing a neurally-inspired cognitive model of decisions where choice options must be self-generated, referred to as "internal-menu choices" (IMC). Specifically, we do so by building upon two well-established principles from the cognitive and neuroscientific literature on memory retrieval and valuation. First, retrieval of semantic knowledge is a probabilistic process governed by the associative principle. Second, given a set of options, choice is governed by subjective preferences over the options, commonly referred to as utility or value. In a series of experiments (N = 3,067) across a diverse array of real-world goods, we show that our computational model was able to make highly precise and accurate predictions (mean out-of-sample R^2 = 0.94) of the likelihood that each option was chosen in IMC by incorporating the extent to which it benefitted from a failure to retrieve some other option, or was passed over due to successful retrieval of another option. Furthermore, using fMRI (N = 32) we confirmed our core assumption regarding the engagement of semantic memory retrieval (lateral prefrontal cortex) and valuation (ventromedial PFC and posterior cingulate cortex) processes in IMC. In addition, functional connectivity analysis showed a significant psychophysiological interaction between vmPFC and left anterolateral prefrontal cortex in IMC, consistent with the idea that the valuation system during IMC relies on communication with retrieval system. In contrast, when the choice was based on an externally-provided set of options, vmPFC showed enhanced functional connectivity with fusiform gyrus, consistent with the interaction with visual processing during menu-reading. Together these

results advance our understanding on the nature of memory retrieval mechanisms underlying adaptive decision-making in the real-world, and can be viewed as a continuation of past work that emphasizes the "boundedly rational" nature of human decision-making. Better understanding retrieval-based differences in decision making within individuals, and in patients for whom these processes may be disrupted, such as in Alzheimer's disease, may represent a first step toward helping them to optimize their decision quality as they define it.

Session V: Dynamic decision making

Individual heterogeneity in foraging behavior relates to drug addiction and midbrain dopamine function

Candace Raio[1], Sara Constantino[2], Kathryn Biernacki[3], Darla Bonagura[3], Joany Xue[3], Kenneth Wengler[4], Guillermo Horga[4], Anna Konova[3]
[1]New York University, [2]Princeton University, [3]Rutgers University?New Brunswick, [4]Columbia University

OBJECTIVE: A prominent feature of addiction is the tendency to exploit a previously rewarding resource despite its diminishing returns. Such behavior is aptly captured in animal foraging models that have recently been extended to humans. Catecholaminergic systems are thought to underlie such behavior, but a precise empirical account of this is lacking in humans. METHOD: We recruited 21 individuals with opioid use disorder (OUD) and 21 socio-demographically matched control participants (25-70 y of age, 43% women). OUD participants were drawn from a larger longitudinal study on heroin relapse in treatment-seeking individuals. Participants completed a foraging task, during which they made sequential decisions between "harvesting" a depleting resource for monetary rewards or incurring a cost to "travel" to a replenished resource. To more directly assess catecholaminergic contributions to foraging behavior, we further acquired high-resolution (<0.7 mm2 in-plane) neuromelanin-sensitive MRI scans, a non-invasive technique used to reliably probe long-term dopamine and norepinephrine function, in a subset (n=27) of OUD and control participants. Our imaging protocol was optimized to separately localize dopaminergic nuclei (substantia nigra, ventral tegmental area, SN/VTA) and the noradrenergic locus coeruleus (LC-NE), both of which have been theoretically linked to foraging behavior and are thought to underlie changes in neural circuits implicated in addiction RESULTS: Our foraging task results replicate previous reports that healthy control participants over-harvest relative to the optimal policy, especially as travel times to new patches increase (travel time effect: p=0.009). Interestingly, OUD participants tended to over-harvest even more than matched controls. This tendency to over-harvest in OUD was coupled with a marked insensitivity to travel times (travel time effect: p=0.79; travel time X diagnosis interaction: F=3.10, p=0.08). These group differences held when controlling for age, sex and cognitive variables, and over-harvesting scaled with increased years of opioid use (OUD: r=0.47,

p=0.048). Our imaging analysis revealed a dissociation whereby, across participants, over-harvesting was associated with lower neuromelanin signal contrast in dopaminergic nuclei (SN/VTA, r=0.41, p=0.03, and separately within VTA, r=0.38), but not in LC (p=0.55). CONCLUSION: Our findings suggest that individual differences in foraging behavior are related to interindividual variability in dopaminergic--but not noradrenergic--circuit function that informs reward rates in dynamic decision environments and may serve as a marker for maladaptive reward-seeking behavior.

Foraging behavior adjusts to multiple scales of context

Nora Harhen[1], Aaron Bornstein[1]
[1]University of California, Irvine

Objective: Most real world environments exhibit temporal variability at multiple levels. The weather may oscillate between rainy and sunny day to day, gradually heat up and cool down season by season, and more slowly shift over decades. Human memory is structured to reflect these various scales of temporal dependencies (Collin et al. 2015). Optimal foragers should consider information across these different timescales when estimating the overall quality of the environment. However, it is not known whether this reference point adapts to context beyond the most recent experiences (Zhang et al 2015; Fougnie et al 2015). Temporally extended contexts are updated more slowly allowing for the integration of even very distant experiences into the representation of the environment's broader structure (Chien & Honey, 2020). Here, we provide initial evidence in favor of multiscale reference points in foraging, by showing that estimation of current patch quality is shaped by both recent and distant experience. Method: We investigated how humans learn in a serial stay/switch foraging task (Constantino & Daw 2015) in which participants decide between staying to harvest a depleting patch of resources or incurring a time delay to switch to a replenished patch. The foraging environment consisted of three patch types of different quality. Patches of the same type were encountered in "clusters", with occasional transitions to a patch cluster of a different type. Participants (N=84) experienced clusters in one of three sequences varying along two dimensions: initial patch type (poor, rich) and extremity of transitions (extreme, graded). Participants were not told that patches varied in quality, nor when transitions would occur, requiring them to infer this information from experience alone. Results: The degree to which participants overharvested, relative to Marginal Value Theorem (MVT; Charnov 1976), was modulated by sequence type ($F(2,579)=13.12$, $p<0.0001$). This broader context effect interacted with local context effects. The extent of overharvesting was influenced by both the current cluster's quality and the preceding cluster's quality, but the direction of influence also depended on the overall sequence of clusters encountered ($F(8,579)=4.53$, $p<0.0001$). Conclusion: These results are consistent with participants evaluating the quality of the current cluster with respect to both recent and distant experience. Our results provide evidence that

while foraging humans track multiple levels of context within the environment producing a dynamic reference point that contrasts with the static threshold prescribed by Marginal Value Theorem.

Assessing trial-level sequential effects in a non-sequential risk task using cognitive models and Bayesian methods

Nidhi Banavar[1], Michael Lee[1]
[1]University of California - Irvine

Objective: Gambling tasks with no explicit sequential structure often assume that data collected in sequence can be treated as independently acquired, as long as participants are instructed to treat each decision as if it were the only one that counts. Is such instruction enough to minimize spillover effects of trials past, and can they be assessed in the absence of feedback? If so, how do they impact inference on behavioral parameters of interest? Here we address these questions using a simple, flexible generative model with psychologically interpretable parameters. Methods: Subjects (n = 56) participated in a gambling task with no explicit sequential dependencies, where they chose between two gambles and were instructed to select the gamble that maximized/minimized reward/loss. Choices were made in both the gain and loss domain (not intermixed). We developed a hierarchical latent-mixture psychometric Bayesian model to assess sequential effects within each condition. The latent-mixture allows for a model-based classification of whether participants follow the normative decision rule (maximize Expected Value [EV]) or are contaminants at the trial, condition or individual level. The hierarchical extension allows the model to infer the parameters characterizing the psychometric function (α "shift" and β "scale"). We modeled potential sequential effects on choice probability (i.e. p(choose left gamble on trial t | left on trial t - 1)) and on the psychometric shift and scale parameters themselves. Results: We find strong evidence of a sequential effect on the scale parameter β in both domains (Bayes Factor [BF] >10 gain; BF > 100 loss): on trials following an "easy" decision (i.e. large difference in EV of the two gambles), the psychometric curve is flatter, and the perceived discriminability between the two present choices is poorer - subjects are less sensitive to differences in EV. Otherwise, we largely find evidence for the absence of a sequential effect. Conclusions: Results suggest that common measures of choice characteristics in ostensibly non-sequential tasks - risk elicitation, discount factor - may indeed be contaminated by presentation order. We provide a straightforward generative method for examining - or excluding - this possibility.

Adaptive arbitration between strategies during human social learning: emulation, imitation and direct experience.

Caroline Charpentier[1], Kiyohito Iigaya[1], John O'Doherty[1]
[1]California Institute of Technology

To navigate our complex social world, it is crucial for people to deploy multiple learning strategies, such as learning from direct experience as well as from social cues. In observational learning (OL), organisms learn from observing the behavior of others, employing least two distinct strategies. Imitation involves repeating other agents' previous actions, while emulation proceeds from inferring their goals and intentions. In experiential learning (EL), organisms learn from directly experiencing the outcomes of their actions. Despite the prevalence of OL and EL in humans and other social animals, a fundamental question remains unaddressed: how is control over behavior assigned to one strategy over another depending on the environment? Here we developed a novel computational model in which arbitration between learning strategies is determined by their predictive reliability, such that control over behavior is adaptively weighted toward the strategy with the most reliable (less uncertain) prediction. We tested this theory in two behavioral tasks: an OL task in which learning can only occur through imitation or emulation (Study 1, three independent samples, total N=230), and an OL-EL task in which participants can learn both from observation and from experience (Study 2, two independent samples, total N=100). Each task design manipulated uncertainty over time and was optimized to produce dissociable effects on the reliability of each strategy. In both studies, behavior manifested patterns consistent with the deployment of the expected strategies (imitation and emulation in Study 1; OL and EL in Study 2) and flexible uncertainty-driven changes between them. Computational modelling revealed that behavior was best explained by an arbitration model, rather than by individual strategy models. Arbitration was driven by the reliability of emulation in Study 1, with imitation deployed when emulation reliability is low, and by the reliability of OL and EL, implemented as separate weights, in Study 2. Being replicated across multiple samples, these findings illuminate a domain-general computational implementation of control over behavior during human social learning. The dynamic and adaptive allocation of control to the most reliable strategy can be efficiently utilized to produce sophisticated social behavior.

Session VI: Efficient coding and representation

Efficient Coding and Risky Choice

Cary Frydman[1], Lawrence Jin[2]
[1]USC Marshall School of Business, [2]Caltech

We experimentally test a theory of risky choice in which the perception of a lottery payoff is noisy due to information processing constraints in the brain. We model perception using the principle of efficient coding, which implies that perception is most accurate for those payoffs

that are expected to occur most frequently. Across two pre-registered laboratory experiments with N=190 subjects and 234,000 decisions, we manipulate the distribution from which payoffs in the choice set are drawn. In our first experiment, we find that risk taking is more sensitive to payoffs that are presented more frequently. These results provide evidence that the principle of efficient coding extends into the domain of risky choice between monetary lotteries. In a follow-up task, we incentivize the same group of subjects to classify which of two symbolic numbers is larger. Subjects exhibit higher accuracy and faster response times for numbers they have observed more frequently. We then estimate a computational model of efficient coding and test for a correlation between structural parameters across tasks. We find that precision of coding in risky choice is correlated with the precision of coding in the number classification task. In our second experiment, we manipulate the payoff distribution so that efficient coding induces the decision maker's perceived value function to switch from concave to convex. We find that demand for risk is significantly higher when efficient coding induces a convex value function. These results also provide novel evidence supporting decision-by-sampling models. Together, our experimental results suggest that risk taking depends systematically on the payoff distribution to which the decision maker's perceptual system has recently adapted. More broadly, we provide novel evidence of the importance of imprecise and efficient coding in economic decision-making.

Efficient learning of statistical regularities in higher-level cognition

Rafael Polania[1], Michael Woodford, Joseph Heng[1]
[1]ETH

The precision of human decisions is limited by both processing noise and basing decisions on finite information. But what determines the degree of such imprecision? It has been suggested that the rules guiding behavior are not arbitrary, but follow fundamental principles of acquiring information from environmental regularities in order to make the best decisions. Moreover, these principles should incorporate strategies of information coding in ways that minimize the costs of inaccurate decisions given biological constraints on information acquisition, an idea known as efficient coding. While early applications of efficient coding theory have primarily been to early stages of sensory processing, it is worth considering whether similar principles may also shape the structure of internal representations of higher-level concepts, such as the perceptions of value that underlie economic decision making. Here we develop an efficient coding framework for higher-level cognitive processes in which information is represented by a finite number of discrete samples. We characterize the sampling process that maximizes perceptual accuracy or fitness under the often-adopted assumption that full adaptation to an environmental distribution is possible, and show how the optimal process differs when detailed information about the current contextual distribution is costly. Here, we report new experimental data on numerosity discrimination by human participants, where we find that our

data are most consistent with an efficient coding theory for which the performance measure is the frequency of correct comparative judgments, and where people economize on the costs associated to learn about the statistics of the environment. Thus, understanding decision behavior requires that we account for biological restrictions on information coding, and our results challenge the often-adopted assumption of precise prior knowledge in higher-level decision systems.

In what environments is divisive normalization an efficient computation?

Stefan Bucher[1], Adam Brandenburger[1]
[1]New York University

Divisive normalization (DN) has been argued to be a canonical computation solving the brain's problem of representing a potentially unbounded stimulus via biophysically feasible bounded firing rates. Originally discovered in the primary visual cortex, DN has since been observed in many other domains including the representation of subjective values. DN is often considered to be an implementation of the efficient coding principle. For example, it has been shown to reduce redundancy in natural stimulus statistics. However, characterizing the environments for which DN is provably efficient has been an open problem. We close this gap with a theoretical result that specifies precisely which stimulus distributions are encoded efficiently by DN. Specifically, we show that, in a low-noise environment, DN maximizes the mutual information between an n-dimensional stimulus and its representation if and only if the distribution of the stimulus in the environment is a multivariate log-logistic distribution. Divisive normalization transforms this particular distribution (and no other) into a uniform distribution, which amounts to a multivariate analog of histogram equalization. We extend our result to allow for an arbitrary metabolic cost of the representation, and show how this impacts the associated stimulus distribution. Interestingly, for some parameter values the log-logistic distribution is similar to the power-law frequency distributions found in natural images, which have played an important role in sensory physiology and empirical studies of DN. We speculate that DN is, in fact, an efficient adaptation to such stimulus distributions occurring in the natural world. Our result yields predictions on how the parameters of the DN function should be tuned to the parameters of the stimulus distribution in the environment. Optimally, the normalization constant should be set to the inverse of the scale parameter, while the exponent modeling the rectification of inputs should depend on the shape parameter of the log-logistic distribution being represented. Our methods can be adapted to other efficiency criteria than maximization of mutual information and, therefore, open the door to further theoretical investigation of the properties of the canonical DN computation.

How representations adapt: the role of structure in model-based planning

Neil Garrett[1], Leonie Glitz[1], Keno Juchems[1], Christopher Summerfield[1]
[1]University of Oxford

Objective: A key aspect to maximising utility involves maintaining an accurate model of an environment's one-step dynamics: which actions take us to which locations? In a world with many actions and states, mapping actions to states with the right level of complexity is critical. Overly complex representations are computationally expensive and provide little opportunity for learning to be generalised to new situations. But representations that are highly compressed risk misapplying updates to unrelated action contingencies leading to poor planning and inaccurate choices. We demonstrate how humans flexibly adapt the task representations they use during economic choice thereby taking advantage of environments that share common structure whilst protecting representations that are environment specific. Methods: We developed the "Heist Task", a learning task in which participants (n=29) needed to track trial by trial fluctuations in state-action contingencies. In dependent blocks, participants made choices in two separate environments, but the underlying transition functions were the same in each. In independent blocks, transition functions were independent. We used computational modelling in conjunction with fMRI to investigate the extent to which the representational architecture used to map actions to states adjusted between these two conditions. Results: First, using a computational learning model, we demonstrate that state representations recruited at the time of choice to compute value, differed between blocks. A lower dimension representation, which enabled information to be integrated between different environments, was favoured in dependent relative to independent blocks. Second, after identifying the involvement of entorhinal cortex and the amygdala in signalling changes to state-action contingencies, we used RSA to compare the similarity of BOLD responses in these brain regions. This revealed a striking similarity in dependent blocks whereby similar state-action contingencies exhibited correlated patterns of BOLD activity between different environments, suggesting the use of a state-action mapping common to both. This similarity was notably absent in the independent condition, consistent with the use of separable state-action mappings, exclusive to each environment, that prevented interference from one to another. Conclusion: Humans are adept at planning, even in environments that are unfamiliar. Our findings propose a computational and neural account of how this occurs. By reducing the complexity of task state representations when environments share common structure, state-action mappings learned in one environment can be used to plan in others.

SYMPOSIUM ABSTRACTS

Symposium 1: Advances in neuroforecasting: forecasting consumer and firm choice using neural data

Using neural data to improve forecasts of market-level behavior

Alexander Genevsky[1], Lester Tong[2], Brian Knutson[2]
[1]Erasmus University, [2]Stanford University

Consequential decisions in business and public policy are often based on forecasts of population-level behavior. Improvements in our ability to accurately predict aggregate-level responses have the potential to maximize the likelihood of success and minimize waste of valuable time and resources. Although a growing literature demonstrating the capacity of neural data to forecast aggregate level behavior (see Knutson and Genevsky 2018), no research has yet explored the mechanisms that account for these out-of-sample predictions nor their optimal conditions. In this paper, we examine the role of generalizable and idiosyncratic choice processes in the prediction of aggregate behavior by exploring how forecasts based on behavioral and neural measures are differentially impacted by the representativeness of the study sample to the market population. Participants in the scanner made incentive compatible decisions regarding projects from a popular online crowdfunding platform. In phase two, a large internet sample (n= 3000) was asked to make similar funding choices regarding the same projects. The laboratory sample's behavior and neural activity were then used to forecast preferences in two markets - one representative of the sample and one less representative. We hypothesized that while in both behavioral and neural measures would be predictive in the well represented market, we expect to observe a steeper decline in accuracy for behavioral prediction and a relatively stable prediction rate for neural measures in the less well represented market. We find that when compared to behavioral measures, neural predictions are less impacted by the representativeness of the sample. Regression models including both behavioral and neural predictors of preference in the representative market indicated that both lab behavior (coef. = .319, SE = .151, $p < .05$) and neural activity (coef. = .613, SE = .211, $p < .01$) accounted for significant and independent variance. However, in the non-representative market, only neural activity was significantly associated with the market preferences (coef. = .656, SE = .217, $p < .01$). These data suggest that neural activity may be a more generalizable index of preference across individuals than self-report measures or observed behavior. In fact, for non-representative samples, only neural activity was a significant predictor of market preference.

Predictive Power of fMRI, Survey and Sales Data - A Data Fusion Approach using Machine Learning to Predict Sales

Anita Tusche[1], Marton Varga[2], Paulo Albuquerque[2], Nadine Gier[3], Bernd Weber[1], Hilke Plassman[2]

[1]Queen's University, [2]INSEAD, [3]Heinrich Heine University

Accurately forecasting the market-level success for new product launches-where no historic sales information is available and failure rates are often high-presents a major challenge. Here we address three questions: can we predict market-level success of newly launched products by jointly using market, survey, and brain data? What are the average contributions of each information source for forecasting? Can brain data improve predictions of market-level purchasing behavior? We used a random forest algorithm to forecast market-level success of 19 food products recently launched by the biggest German supermarket chain (test data). Market-level success was captured in sales and the fraction of adopting local retailers of each product. Predictions were based on three types of data of 37 previously launched substitute products (training data) chosen based on the retailer's sales-based substitute index: (1) sales data of substitutes (N=114345 data points; i.e., historic sales); (2) survey data assessing product attitudes (representative online sample, N=1451; ~300 evaluations per product; i.e., product desirability); (3) brain data (fMRI) in three areas of the brain's valuation system obtained in a purchasing task (convenience sample, N=44; average estimated brain activity in three a-priori defined regions of interest: nucleus accumbens, ventromedial prefrontal cortex, anterior insula). We compared forecast performance (Root Mean Squared Error) for different models that systematically varied the data considered for the out-of-sample prediction of market-level success. Taking advantage of all three data types (full model) reduced prediction error by 39% over the baseline model. Sales data contribute most to the prediction of sales (53%), followed by survey data (28%) and fMRI data (19%). Model comparisons revealed that the addition of brain data - obtained in a small, non-representative subject sample - improved predictions of the market-level success of newly launched products in 65% of cases. We show that combining a retailer's sales data about competitors with information from traditional surveys and fMRI data can dramatically improve predictions of market-level success of newly launched products. Gauging the added value of each information source reveals that while historic sales and survey data contributed heavily to predictions of future sales, brain data improved forecasts of market-level success for the majority of cases. Our results illustrate what managers can learn from brain data to forecast market preferences, supporting decision whether to launch a product, take it off the market, effectively plan operations and adjust marketing policies.

Population news sharing is reflected in distributed reward-related brain activity

Bruce Doré[1], Christin Scholz[2], Elisa Baek[3], Emily Falk[4]
[1]McGill, [2]University of Amsterdam, [3]UCLA, [4]University of Pennsylvania

Information that diffuses widely in the media environment can influence the behavior of individuals and shape broader directions of societal change. Previous studies have shown that information in targeted brain regions can be used to predict this diffusion of information. However, by focusing on these targeted regions, the approach in previous work has discarded information from the majority of the brain. Current neurobiological models posit that the functional neuroanatomy of reward value extends beyond the core striatal and ventromedial prefrontal (vmPFC) regions that have received empirical attention in previous population prediction studies. These models propose that diverse brain systems interact to rapidly propagate reward-related information throughout the brain, generating a distributed value representation that directs cognition and behavior in a multi-faceted manner. However, it is unclear whether distributed brain representations of reward value hold information that can be used to predict the impact of messaging, beyond responses in core striatal and vmPFC regions. We sought to address this gap in knowledge with two neuroimaging studies that quantified functional brain responses to New York Times health news articles and used these responses to predict sharing of these articles in the broader population of readers. In particular, we addressed two specific questions. First, does expression of a meta-analytically defined reward valuation-related pattern predict population-level information sharing? Second, does expression of this pattern improve prediction of population-level information sharing beyond what can be predicted from activity within reward-related brain regions and self-reports? Our results indicated that expression of a distributed pattern of brain activity meta-analytically associated with reward valuation substantially improved accuracy in predicting population sharing of the news articles, beyond previously identified brain and self-report predictors. Further, the predictive efficacy of the pattern was not reducible to patterns of activity within core brain reward regions but rather depended on larger-scale patterns of activity distributed widely across cortical, subcortical, and brainstem systems. These findings highlight the advantages of using whole-brain patterns in addition to previously used region-specific and self-report predictors of information sharing. Overall, this work contributes to our growing ability to use neural data to forecast out-of-sample outcomes and augments our understanding of the brain mechanisms that underlie how information diffuses (or fails to diffuse) across a population of individuals.

Neuroforecasting aggregate choice in online dating: Predicting aggregate choices from small samples using neural and behavioral measures

Steven Shaw[1], Vinod Venkatraman[2], Carolyn Yoon[1]
[1]University of Michigan, [2]Temple University

The ability to predict aggregate, market level choices from a small sample of individuals can provide tremendous value in several domains (e.g., product success, political elections). In this paper, we advance our understanding of decision-making and choice by assessing behavioral and neural approaches to forecasting aggregate choices in an online dating context, both within a small sample (n < 50) and an independent simulated market sample (n > 600). In accordance with the affective-integration-motivation (AIM) framework, we propose that affective neural components of individual choice are most useful for aggregate forecasting, whereas neural components associated with the integration of information are most useful for individual-level prediction. We developed 36 standardized dating profiles with orthogonalized dimensions, using a factorial design: attractiveness (high, medium, low), age (19-23, 24-28), facial expression (neutral, smiling), and profile description (hobbies/likes, SES/occupation, and personality traits). Respondents made a binary 'like' or 'pass' choice for each profile, and rated each profile on attractiveness, career prospects, likability of personality, and likelihood that the individual in the profile will 'like' them back (i.e., choice decomposition variables). Training models with aggregate in-lab sample levels of choice decomposition variables to predict the aggregate choice likelihood of randomly selected holdout profiles (75/25 train/test; 50 iterations) resulted in correct choice likelihood prediction (+/- 5% of actual choice likelihood) 66.3% of the time within sample, and 38.3% of the time for the market population. Additionally, we observe a dissociation between the choice decomposition variables that are predictive of within sample versus market level aggregate choice; for within sample choice, aggregate sample-level attractiveness, career prospects, and likeback variables significantly predict aggregate choices, while for market level aggregate choice, aggregate sample-level attractiveness alone is predictive. In accordance with the AIM model, we hypothesize: H1: Affective neural components, captured by activity in the NAcc, will be most highly correlated with perceived attractiveness and facial expression, and more likely to forecast aggregate choice. H2: Neural activity associated with integration, captured by activity in the vmPFC, will be most highly correlated with perceptions of career prospects and personality and will be most useful for predicting choices at the individual-level. At this time, neuroimaging data collection on a small participant sample has been postponed due to the COVID-19 pandemic.

Symposium 2: Neuroeconomics meets the digital age

A computational reward learning account of social media engagement

Björn Lindström[1], Martin Bellander[2], David Schultner[1], Allen Chang, Philippe Tobler[3], David Amodio[4]

[1]University of Amsterdam, [2]Karolinska Instiutet, [3]University of Zurich, [4]New York University

Social media has become a modern arena for human life, with billions of daily users worldwide. The intense popularity of social media is often attributed to a psychological need for social rewards ("likes"), portraying the online world as a "Skinner Box" for the modern human. Yet despite such portrayals, empirical evidence for social media engagement as reward-based behavior remains scant. We applied a computational approach to directly test whether reward learning mechanisms contribute to social media behavior. We analyzed over one million posts from over 4,000 individuals on multiple social media platforms, using computational models based on reinforcement learning theory originally developed for animal free-operant behavior. Our results consistently show that human behavior on social media conforms qualitatively and quantitatively to the principles of reward learning. Specifically, users posted with a higher rate when the environmental reward rate was higher, relative to lower. Results further reveal meaningful individual differences in social reward learning on social media, explained in part by variability in users' tendency for social comparison. An online experiment (n = 176), based on a simulated social-media environment, verified that social rewards causally influence behavior as posited by our computational account. Together, these findings support a reward learning account of social media engagement and offer new insights into this emergent mode of modern human behavior.

Using Neuroscience to Causally Affect Information Sharing

Christin Scholz[1], Elisa Baek[2], Emily Falk[3]
[1]University of Amsterdam, [2]UCLA, [3]University of Pennsylvania

The frequency and manner in which information is shared socially, e.g. on social media, both the reach and impact of content on target populations. Important information which could positively influence society, like verified health content, may not naturally have characteristics that facilitate virality (e.g. emotionality, novelty), and, thus, frequently underperforms in today's information environment. Rather than focusing on features of viral content, neuroscientists have identified neural and inferred psychological mechanisms that precede sharing, highlighting the involvement of brain regions that are meta-analytically associated with self-related and social processing. There is hope that identified mechanisms may inspire novel interventions that affect sharing by encouraging sharing-relevant thoughts. To test the validity of previous reverse inferences and develop a practical, causal sharing intervention, we ran a pre-registered (https://osf.io/n9vpz/) fMRI experiment (N=53) in which a group of communicators considered sharing health news articles with receivers. In a within-subject manipulation, communicators were asked to use articles to "Describe Yourself" (DY; encouraging self-related thought), to "Help Somebody" (HS; encouraging social thought) or to objectively "Spread Information" (SI; control). Articles were randomly assigned to condition to rule out effects of content features. As expected, both experimental conditions increased brain activity in self-related (DY: M=0.002; CI[0.001, 0.004]; HS: M=0.002, CI[0.001, 0.003]) and social

regions of interests/ROIs (DY: M=0.004, CI[0.001, 0.007]; HS: M=0.003, CI[0.001, 0.005]) relative to SI. The HS(>SI) condition also causally increased sharing intentions (B=0.13, SE=0.06, p=.02). DY(>SI) did not directly but indirectly affect sharing intention through self-related ROI activity (B=0.016; 95% CI [0.004, 0.045]). LIWC scores of communicators sharing texts further revealed pronounced condition differences in the language used to share content (e.g. in social language use, first-person singular, and health-related words). In sum, we confirmed reverse inferences regarding psychological drivers of sharing and causally changed both brain activity in a priori ROIs and sharing intention using a simple intervention. At SNE2020, we will also discuss downstream intervention effects on brain activity and perceptions in a group of 50 receivers who read communicator sharing texts while undergoing fMRI. This study constitutes a practical contribution of neuroscience to a timely problem: How do we support sharing of high-quality health news irrespective of content characteristics in the age of misinformation?

Social learning of moral outrage in online networks

William Brady[1], Killian McLoughlin[1], Tuan Doan[1], Molly Crockett[1]
[1]Yale

Objective: Moral outrage is a core feature of human morality with important consequences for society. In the modern digital era where social media platforms dominate communication, the visibility and social transmission of moral outrage has increased with major consequences for individuals and organizations. This research aimed to understand how outrage is amplified in online social networks. Our central proposal was that online outrage is a product of domain general social learning mechanisms that are amplified by the unique design features of social media platforms. Specifically, we examined the streamlined delivery of social feedback and the organization of users into large social networks where norms of expression are readily observable. Methods: In two observational studies, we collected the tweet histories of over 7,000 Twitter users (over 11 million observations) and measured daily moral outrage expression by developing a novel machine learning classifier. Controlling for key confounding effects, we modeled the effect of past social feedback (likes and shares) on current outrage expression to test how users' outrage could be predicted as a function of social feedback history. We also estimated the political ideological extremity of users' social network to test how users in more extreme networks with higher frequencies of outrage expression respond to social feedback compared to users in less extreme networks. Two behavioral experiments created a mock social media environment to directly manipulate social feedback for outrage vs. neutral tweeting, as well as frequency of outrage in their "social network" to allow for replication with causal inference. Results: We confirm that outrage expressions on social media increase as a function of positive prediction errors in past social feedback, consistent with principles of reinforcement learning. However, users in ideologically extreme networks where outrage was more prominent were less sensitive to social feedback even though they expressed more outrage overall.

Behavioral studies confirmed that when outrage is frequent in a network, people are more likely to guide their outrage expression based on observational learning compared to trial and error reinforcement learning (although both impacted outrage tweeting). Conclusions: We provide the first evidence that moral outrage in online networks can be partially explained by trial and error reinforcement learning and observational learning that are amplified by two key features of social media platforms: streamlined delivery of social feedback and organization of users into large social networks with norms of expression readily observable.

Polarized neural responses to political content are associated with biased assimilation of political information and subsequent attitude change

Yuan Chang Leong[1], Janice Chen[2], Robb Willer[3], Jamil Zaki[3]
[1]University of California, Berkeley, [2]Johns Hopkins University, [3]Stanford University

When presented with the identical political content, people tend to interpret information in a manner that supports their prior beliefs. This bias is thought to contribute to increasing attitude polarization. What neural processes drive the divergent processing of political information and what effects do they have on attitude change? To address these questions, we employed a multimethod approach that combines fMRI and semantic analyses of real-world political content. 38 American participants with conservative-leaning or liberal-leaning immigration attitudes watched 24 videos related to immigration policy while undergoing fMRI. These videos were taken from youtube.com and include news clips, campaign ads, and public speeches from prominent politicians. Using inter-subject correlation analyses, we searched for evidence of "neural polarization": activity in the brain that diverges between people who hold liberal versus conservative political attitudes when watching the videos. We observed neural polarization only in the dorsomedial prefrontal cortex (DMPFC), a brain region previously shown to track the interpretation of narrative content. To assess the content features associated with biased processing, we broke down the content of the videos into 50 semantic categories and tested the extent to which each category was associated with neural polarization. Neural polarization in the DMPFC increased during moments in the videos that included risk-related and moral-emotional language, suggesting that these message features are likely to drive divergent interpretations. Furthermore, the degree to which a participant's DMPFC response was similar to that of the average conservative or average liberal participant predicted attitude change towards conservative or liberal positions respectively, suggesting that adopting an interpretation closer to a particular group biased participants towards the positions held by that group. How might neural polarization in the DMPFC arise? One possibility is that inputs to the DMPFC are modulated by one's political attitudes. Consistent with this hypothesis, we found that political differences were associated with divergent frontostriatal connectivity. Together, our results provide a neural account of how biased processing in the brain gives rise

to divergent interpretations of political information and subsequent attitude polarization. These findings thus shed new light on the neural processes that influence political choice.

POSTER ABSTRACTS

F-1: Underlying mechanisms of purchase decision-making in e-commerce

Nadia Morozova[1]
[1]Warwick Business School

The main objective of this study was to achieve an in-depth understanding of the nature and underlying mechanisms of the decision-making process that leads to planned, unplanned and impulse purchases. Thus, to identify shoppers' emotional engagement and cognitive involvement in the decision-making process, we combined physiological data from E4 wristbands (electrodermal activity (EDA), heart rate (HR) and skin temperature), and mouse tracking data collected in the course of a shopping trip. To determine the type of purchase we conducted a systematic literature review and based on its results developed and tested a scale. Moreover, we created research stimuli (e-commerce site) with the use of the Design Thinking approach. The main research experiment included 71 respondents. Linear Mixed Model analysis highlighted a significant relationship ($p<0.05$) between shoppers' physiological responses and the type of purchases made in the e-store at the motivation stage of purchase decision-making (last 20 sec of product exposure). In addition, the analysis shows that HR activation is significantly higher ($p<0.05$) for impulse purchases in comparison with unplanned and planned purchases, which corresponds to EDA trends. ANOVA analysis of the mouse tracking data indicates that the unplanned purchase decision-making process takes significantly less time ($p<0.05$) in comparison with impulse and planned purchases. This result goes in line with significantly higher ($p<o.05$) HR activation at the attention stage of the purchase decision-making process (first 5 sec of product exposure) in the case of unplanned purchases. Considering different site triggers that shoppers use during their purchase decision-making, the Chi-Square test highlights a significant association ($p<0.05$) between site triggers and types of purchases made on the site. Further crosstabulation analysis revealed that shoppers use significantly more ($p<0.05$) Specific site triggers and significantly less ($p<0.05$) General site triggers (e.g. list of products and product page) for impulse purchase decision-making than for unplanned purchases. Hence, we can conclude that shoppers' attention towards products upon first exposure plays a more critical role for unplanned purchase decision-making, while in the case of impulse purchases emotional engagement at the motivation stage of the decision-making process is of higher importance. Moreover, shoppers follow different paths to purchase while making planned, unplanned and impulse purchases. Key research recommendations have been implemented in the business strategies of such companies as Pinterest, Google, RB, L'Oreal and MediaCom.

F-2: Neural signature of liking during video watching

Hang Yee Chan[1], Ale Smidts[1], Maarten Boksem[1]
[1]Rotterdam School of Management

A main goal of consumer neuroscience is to predict consumer preference based on brain activity. We used machine learning algorithm (GraphNet) to identify neural signature of liking of TV commercials. The neural signature predicted self-report liking in a cross-validation sample (N = 25) watching 20 commercials, and a separate larger sample (N = 60) using a different set of 35 commercials. Predictions based on whole-brain neural signature achieved 80% classification accuracy in video liking in the testing data, and outperformed traditional region-of-interest analysis with brain areas known to be associated with TV commercial liking. We further investigated the temporal dynamics and found the expression of neural signature of liking emerged within the first 5s of watching TV commercials. This study provides a basis for new brain-based method to predict consumers preference in dynamic experience such as video watching.

F-3: The Effects of Psychological Perceptions on Brand Names Associated to the Coronavirus Pandemic

Yu-Wen (Diana) Shih[1], Aline Simonetti[2], Enrique Bigné[2]
[1] University of Bonn, [2]University of Valencia

The horn effect bias is the perception that negative traits are connected without having an intrinsic connection (Nisbett & Wilson, 1977; Sigall & Ostrove, 1975). For example, the genetically modified food label can cause a horn effect on consumers' calorie estimates (Burton et al., 2015). However, it is unclear how the horn effect could influence consumers' decision-making process in brands in the current COVID-19 pandemic, more specifically, for a brand name that evokes the same semantic association with the coronavirus. We conducted a survey study through Clickworker between the 4th of May to the 9th of May 2020 in the U.S. (n = 414). Three alcoholic brands were chosen in this study - Corona Extra beer, which has a direct association with the coronavirus (i.e., Corona); Modelo Especial beer, which has similar brand features with Corona Extra beer, was chosen as a control brand; Finally, Virus Vodka was chosen as an indirect association with the coronavirus (i.e., Virus). Each participant answered questions regarding the Big Five Personality scale, level of blaming the brand in attribution to the pandemic situation, brand familiarity, overall feeling toward the brands, consumption frequency of the chosen brands before and during the pandemic, and psychological influences of the pandemic. Our analyses focused on participants' risk-perception of the COVID-19 pandemic, whether they blamed the brands for the occurrence of the pandemic and how it affects their frequency of consumption. The findings suggest that, on average, participants perceived the pandemic as a high-risk situation and reported a medium level of perceived anxiety and severity. Relative to pre-pandemic moments, participants' consumption frequency of the three alcoholic brands did not differ during the pandemic. Moreover, neither the overall

feeling toward the brands, nor risk-perception, anxiety, and severity levels was significantly associated with the brands and the blame attribute. The correlation between the overall brand feeling and the pandemic-related metrics was weak and positive for Corona Extra and Modelo Especial, and null for Virus vodka. The strength and direction of the relationships were similar for the three brands, indicating a null effect of participants exhibiting the horn effect bias for these alcoholic brands due to the COVID-19 pandemic. Overall, consumers perceived feeling toward the brands was not associated with the negative impacts of the pandemic and their consumption frequency was not affected. Despite the conspiracy theory, as well as a spike in searches for Corona and "beer virus" on Google Trends (Deese, 2020), this study showed no evidence that the negativ

F-4: Using mobile eye-tracking to capture the effects of choice set size on information processing during purchase decisions in the field

Alexandre Filipowicz[1], Laura Zaneski[1], M. Kathleen Caulfield[2], Quentin Andre[3], Eric Singler[4], Hilke Plassmann[5], Joseph Kable[1]
[1]University of Pennsylvania, [2]University of Pittsburg, [3]Erasmus University, [4]BVA Nudge Unit, [5]INSEAD

Laboratory studies often find that choice behavior is influenced by the number of available options: as the number of options increases, people narrow their focus on a subset of the options available. We tested whether this general finding applies in more naturalistic shopping settings. In a lab-in-the-field study, we measured eye gaze with a mobile eye-tracker while participants moved through a professional mock grocery store and selected items for purchase from three categories: bread, chocolate and sugar. Across participants (N=224, 162 female, mean age=35 years), we manipulated whether the number items available in each category contained 3, 6 or 12 different options. Looking behavior was measured during a participant's 'choice period' for each item category, defined as the period between a participant's first fixation on an in-category item and the point at which they make their choice. We found that, as the number of options increases, shoppers become less likely to look at all available options (mean probability of looking at all items in the 3, 6, and 12 item conditions: 0.96,0.89,0.73 respectively; all pairwise comparisons ps<.001), preferring instead to look at items that match the brand they eventually choose and skipping over other items (p=.014). This effect was most prominent for item categories in which shoppers expressed that brand mattered for their choice (bread and chocolate; ps<.003), but not for categories in which brand had little influence (sugar; p=0.090). Additionally, this bias towards chosen brand starts earlier in the choice period as the number of choices increases, with shoppers becoming increasingly likely to first fixate on the brands they eventually chose (p<.001). This reduction-by-attribute strategy mirrors what has been found in laboratory studies: as the number of options increases, people stop evaluating every option and instead focus on a subset of options. In our case, as the number of option increases, shoppers become more likely to use brand preference to guide their

decisions, particularly when brand matters for their choice. Our results demonstrate the usefulness of mobile eye-tracking for testing hypotheses about information processing in more naturalistic settings. They also have implications for marketers, proposing that when shelf space is scarce, branding is less important and customers might be more likely to churn.

F-6: Heuristics and auditory stimuli affect visual information processing in financial decisions

Maria Gabriella Ceravolo[1], Vincenzo Farina[2], Lucrezia Fattobene[3], Lucia Leonelli[2], GianMario Raggetti[4]
[1]Politecnica delle Marche University, [2]University of Rome Tor Vergata, [3]LUM Jean Monnet University, [4]Brain Line

Before opening a bank account, clients must read and sign the Payment Account Fees Information Document (hereinafter, FID), which details product costs. In the attempt to filter the most relevant information from the FID, subjects might anchor their evaluation to a few items, such as the Annual Fee. Moreover, the increase in sensory inputs, especially if conveyed through a different sensory channel (like the auditory one) might challenge the attention capacity, interfering with the information processing and the final decision. In this eye-tracking study, 70 participants (23±2 years; 52 males) were tested. The experiment protocol included 24 visual stimuli representing as many couples of FIDs, displayed on a laptop screen. Participants had to declare which one, for each couple of FIDs, was more attractive. Stimuli presented mild differences between FIDs, so that one was objectively more advantageous than the other. FID couples were matched so that, in half of them, the Annual Fee of the advantageous FID was disadvantageous (i.e. higher than that of the more expensive FID) (dissonant condition). Each participant was exposed to three different environmental conditions (ambient music, buzz, silence), randomly. Six rectangular areas of interest (AOIs), three per FID, were defined: Liquidity, Fixed Cost (FC) Payments, Variable Cost (VC) Payments. Overall, subjects failed to recognize advantageous products in 20% of the cases when the information was dissonant, compared to 10% when it was consonant, supporting the hypothesis of an anchoring bias towards the Annual Fee. Congruently, the AOI Liquidity (displaying the Annual Fee) grabbed the highest attention; moreover, higher Fixation Time (FT) and more revisits to the AOI Liquidity were detected when subjects answered incorrectly. During the music or buzz conditions, higher FT were observed compared with silence, even if this behaviour did not lead to a significant increase in incorrect choices. Financial literacy also modulated attention, as subjects with the lower scores showed a stronger anchoring to the Liquidity. This study confirms the use of heuristics in the visual scan of mandatory disclosure documents. Anchoring behaviour is stronger in subjects with poor financial literacy. The exposure to auditory stimuli, though acting as a distractor, does not significantly affect financial evaluation. Findings are relevant for regulators and bank managers responsible for documents crafting and for disciplining the environmental conditions in which consumers read and sign contracts.

F-7: Neuroforecasting of the consumer choice

Andrew Kislov[1], Anna Shestakova[1], Vadim Ushakov[2], Valeriia Beliaeva[3], Olga Savelo[1], Vasily Klucharev[1]
[1]Institute of Cognitive Neuroscience, Higher School of Economics, [2]National Research Centre "Kurchatov Institute", Moscow, Russia, [3]ETH Zurich

Neuroforecasting refers to the ability to use brain activity of a group of participants to forecast the behaviour of an independent group. A number of the neuroforcasting studies has shown that it is possible to predict a subject's value-based choice using functional magnetic resonance imaging (fMRI) data at the aggregated group level. Such decisions appear to be associated with song popularity(Berns et al., 2012), market-level microlending(Genevsky&Knutson, 2015), crowdfunding outcome(Genevsky et al., 2017), and videos efficiency(Falk&Berkman, 2011; Venkatraman et al., 2015; Tong et al., 2020). Despite the growing popularity of neuroforecasting, no single study has investigated a possibility to forecast aggregated choices of primary rewards. Therefore we made an attempt to (neuro)forecast aggregated choices of various dishes in a popular restaurant. During our fMRI study, thirty participants were exposed to 84 photos of dishes from a menu of the popular restaurant chain. In addition to self-reported preferences, fMRI data was extracted from a priori domain-general and task-specific ROIs including the ventral striatum (VS). We correlated these neural activities with total real one-year sales information provided by the restaurant chain. Activity in the VS to photos of dishes significantly correlated with one-year sales, r=0,24 (p=0,03). Overall, our results confirm prior studies showing that brain activity in the reward system of a relatively small number of participants can forecast the aggregate choice of a larger independent group of individuals. However, reliability of the neuroforecasting critically depends on the Region Of Interest, which should be reward-specific. The study was supported by the International Laboratory of Social Neurobiology HSE,RF Government grant,ag.No.075-15-2019-1930

F-8: Temporal normalization during valuation creates preference reversals

Ursa Bernardic[1], Maël Lebreton[1]
[1]University of Geneva

Objective: Preference reversal refers to a systematic disparity between people's valuation of options and choices between these. Attempts to account and model such preference reversals draw attention across multiple fields (psychology, economy, and neuroeconomics) for more than four decades. Here, we investigate how the distribution of items faced during valuation affects sequential valuation and, ultimately, choice efficiency. To this end, we experimentally manipulated the sequence of options being presented in a valuation task and evaluated how the elicited values could predict subsequent binary choices. Methods: We carried out an

experiment in which subjects (n=41, age: 19-25 y.o., 20 females) performed 1) a value rating task, and 2) a binary decision-making task. In the value-rating phase, participants rated on a scale from 0 to 10 CHF how much they would be willing to pay for a presented item belonging to one of two categories: food items or trinkets from one of the biggest supermarkets in Switzerland. Importantly, unbeknown to participants, the order in which these two categories were presented was manipulated: the items were split into two blocks, separating high and low-value items of this category, based on the actual price of items. In the manipulated category participants evaluated items from low (or high) and afterward from the high (or low) block, while for the non-manipulated category, items were drawn randomly from either low or high block. To evaluate the impact of our manipulation on valuation consistency, the binary decision-making task consisted of 50% within (both items from the same block) and 50% between blocks (each item from a different block) choices for each category. In the end, participants also rated how familiar they were with the presented items. Results: Choice data was analyzed with a generalized linear mixed-effects regression. More precisely, binary preference decisions were fitted with a logistic function and the value rating for the preferred and non-preferred options as dependent variables. Our preliminary results revealed that manipulating the distribution of object values impacts valuation and causes preference reversals. As such, the value difference of the options predict binary choices for the manipulated, but much weaker for the non-manipulated category. Further analyses suggest that temporal normalization in valuation is domain-specific and not generic. Conclusions: Taken together, our results both shed light on the new experimental mechanism for studying temporal normalization and increase understanding of preference reversals as a result of temporal normalization mechanisms.

F-9: Loan officers' decision strategies: an exploratory eye-tracking study

Maarten Boksem[1], Lucrezia Fattobene[2], Frank G.H. Hartmann[1], Aljaz Sluga[1]
[1]Rotterdam School of Management, [2]LUM Jean Monnet University

Loan officers' screening of potential borrowers is crucial in the bank lending activity, and with the global financial crisis, the importance of studying loan officers' decision making has become even more clear. Existing research in judgment and decision-making demonstrates that to infer values on some criterion of interest, people sometimes rely on information-intensive strategies that attempt to integrate all available information and other times use information-frugal strategies that ignore some of the information, and that the latter may be more common among expert decision makers. The main goal of the current study is to investigate the decision making strategies loan officers use when making lending decisions. This goal is pursued through two main analyses based on the data from an experiment that included 42 professional loan officers working at a major bank. During the experiment, loan officers were presented with a series of 30 decisions. In each decision, loan officers chose which of the two companies described by a series of financial and non-financial indicators (cues) is more credit-worthy. In

the first analysis we applied a modern Bayesian outcome-based classification approach to model loan officers' decision strategies, featuring three decision strategies: the information-intensive and compensatory Weighted Additive strategy (Payne et al. 1993), the information-frugal and noncompensatory Take The Best strategy (Gigerenzer and Goldstein, 1999) and the Equal Weights strategy (Einhorn and Hogarth, 1975). The second analysis focused on the process of attention allocation over the financial and non-financial information describing the potential borrowing companies and was based on eye-tracking data, to validate the modeling findings about loan officers' use of decision strategies, and to extend them by exploring other features of information processing. We investigated whether loan officers relied on cue-based processing or on alternative-based processing; moreover, the reliance on importance-based or position-based search was tested. Results suggest that the majority of loan officers employed information-intensive strategies, with the majority classified as users of a compensatory strategy. The eye-tracking analysis corroborated the modeling results, and revealed that the visual search of most loan officers was more consistent with either cue-based processing or a mixture of cue-based and alternative-based processing, and also showed that position-based search order predominated, i.e. the process of attention allocation was influenced by the position of the information more than its importance.

F-10: Neuroforecasting political campaign survival

Lester Tong[1], Jon Krosnick[1], Brian Knutson[1]
[1]Stanford

Can the fate of political campaigns be forecast? Some theorists claim that early implicit responses to candidates drive eventual campaign success, while others argue that the course of campaign events determine their eventual success. In this research, we set out to determine whether neural activity in a group of subjects (n=46) could forecast the survival of the political campaigns of 15 candidates during the Democratic primary of 2019-2020, months after the neural data were collected. Consistent with a partial scaling account of neuroforecasting, we predicted that individuals? Nucleus Accumbens (NAcc) and Medial PreFrontal Cortex (MPFC) activity might predict individuals? candidate endorsement, while sample neural activity, but not sample endorsement choices, would forecast candidates? campaign survival months later. We also tested whether neural measures could outperform more conventional forecasts (e.g., from concurrent polls). Results indicated that a combination of NAcc and MPFC activity not only predicted individual endorsement, but also forecast political campaign survival, even when individual behavior did not. Theoretically, these findings support a partial scaling account of neuroforecasting by demonstrating that sampled neural activity (but not behavior) can forecast political as well as market outcomes. Practically, the findings hold implications for enhancing the efficiency and impact of political campaigns by suggesting that candidate selection plays a critical early role in campaign survival.

F-11: Faster processing of gains associated with reduced aversion to loss

Nicolette Sullivan[1]
[1]Duke University

Individuals tend to be more sensitive to losses than gains. The potential for negative outcomes triggers a large range of behavioral anomalies, which are typically accounted for by positing that losses are felt, and therefore weighted, more in decisions. However, this assumes that individuals exhibiting loss-averse behavior hold a true underlying preference for overweighting loss. This paper proposes an alternative framework for loss aversion under which an individual could hold a loss-neutral preference, yet still exhibit loss-averse behavior. Specifically, the speed with which losses and gains are incorporated into the decision process may influence an individual's choices independent of preferences. To test this, computer mouse tracking was used during a gambling task to estimate the time at which gains and losses begin to influence the decision process. Results indicate that although gains enter the decision process first, differences in the time at which losses and gains begin to influence the choice process explained a large degree of individual variance in loss aversion. Faster loss latencies, and slower gain latencies, meant that information about potential losses are available to the decision maker for relatively longer during choice, and lead to greater loss aversion. In addition, loss and gain latencies fluctuated across individuals and trials; in trials with faster gain than loss processing, the gamble was selected. Together, these results suggest that a large degree of the observed overweighting of losses, relative to gains, can be explained by slower latency of a gamble's potential benefit, rather than only an aversion to loss.

F-12: Food prices in the retail food environment modify genetic susceptibility to obesity

Hannah Yang Han[1], Catherine Paquet[2], Laurette Dubé[1], Daiva Nielsen[1]
[1]McGill University, [2]Université Laval

The role of the food environment in obesity risk is unclear, which may be due in part to imprecision in food environment measurements and lack of consideration of neurobehavioral differences in responsivity to environmental food cues. Body mass index (BMI) is known to be heritable, and some variants implicated in genetic risk of obesity are involved in neurobehavioral food cue responsivity. However, whether in-store retail food cue stimuli associate with diet or obesity outcomes according to genetics is unknown. The objective was to assess interactions between everyday exposures in the retail food environment and genetic susceptibility to obesity on diet and adiposity outcomes. Genetic, dietary, lifestyle, and anthropometric data from the CARTaGENE Quebec biobank were linked to provincially representative in-store retail food marketing data by forward sortation area (n=3,718 participants). Diet quality was assessed using the Canadian-adapted Healthy Eating Index. BMI

(kg/m²), waist circumference (cm), and body fat percent were physically measured for all participants. Genetic risk of obesity was assessed with a polygenic risk score (PRS) comprised of 97 genetic variants associated with BMI. Exposure ratios were calculated for in-store display, variety, regular price, and price discount frequency of vegetables in relation to soft drinks, providing an indicator of healthful to unhealthful exposure. Sex-stratified generalized linear models adjusted for neighborhood demographic factors were performed to evaluate associations and gene-environment interactions on diet quality and adiposity outcomes. A significant interaction was observed with regular price, such that the positive association between PRS and waist circumference was strengthened with increasing price of vegetables in relation to soft drinks (estimate: 0.7, 95% confidence interval (CI): 0.1, 1.3). Among men, the in-store display ratio was associated with BMI (-0.3, CI: -0.5, 0.0) and waist circumference (-0.8, CI: -1.5, -0.1) and an interaction was observed for PRS*price discount frequency on BMI (-0.4, CI: -0.7, -0.2), waist circumference (-1.2, CI: -1.9, -0.5), and body fat % (-0.5, CI: -0.9, -0.1). Among women, a significant main effect association was observed between regular price and diet quality (-0.6, CI: -1.1, 0.0) and a PRS*regular price interaction on body fat % was also observed (0.4, CI: 0.0, 0.8). Sex differences are apparent in diet and adiposity outcomes according to in-store retail food environment exposures. These results suggest that prices of low-energy density foods and consideration of sex-specific food cue responsivity are targets to address population obesity rates.

F-13: Utilizing Neural Responses to Prior Trials to Improve Predictions of Willingness-to-Pay

John Clithero[1], Danielle Cosme[2], Bryce Dirks[3], Junaid Merchant[4], Nicole Giuliani[1], Elliot Berkman[1]
[1]University of Oregon, [2]University of Pennsylvania, [3]University of Miami, [4]University of Maryland

Neuroeconomics has made tremendous progress in understanding the computation of subjective value. One application of these advances uses neural data to improve predictions about choice behavior. However, in the neural domain, the existing literature focuses on predicting choice with concurrent brain responses. Here, we expand the prediction exercise to determine the utility in also considering earlier brain responses. We examine this issue using functional magnetic resonance imaging (fMRI) data on a willingness-to-pay (WTP) task and supervised machine learning (SML) methods. Participants (N=94) aged 35-46 (M = 39.2) were recruited as part of a six-month longitudinal intervention study to improve healthy eating habits during middle age. Participants completed an incentivized WTP task involving snack foods. Two sets of brain regions of interest (ROIs) were constructed to use as feature spaces for SML. First, three ROIs from a meta-analysis of subjective value: one for each of ventromedial prefrontal cortex, ventral striatum, and posterior cingulate cortex. Second, 200 ROIs from an established parcellation of human cortex. Prediction performance of combinations of ROIs - for either the current trial only, previous trial only, or both - were compared using SML. We used leave-one-out cross-validation: models were trained on N-1 participants and tested on the remaining one. All features were mean-centered and scaled, and all hyperparameters were

tuned via grid search. For prediction performance we report the average root mean-squared error (RMSE). Our primary objective was to predict an individual's WTP on a single trial. Several findings converge on the idea that fMRI data from the previous trial offer uniquely predictive information regarding a current trial's subjective value. As a baseline, we used the previous trial's WTP as the sole out-of-sample predictor in a linear regression (RMSE = 0.756). A linear regression model adding all ROIs from the previous trial improved prediction (RMSE=0.690). Using the least absolute shrinkage and selection operator (lasso) on the same set of features further improved prediction (RMSE = 0.662). When both the current and previous trial ROIs are considered, lasso performance further increases (RMSE = 0.5738). Importantly, 41% of the features remaining under an optimized lasso were from the previous period. We demonstrate the fMRI data from a previous trial, on its own, provides improved WTP predictions above an easily observed behavioral measure. When data from the current trial are included, the predictive power of the previous trial remains. Both of these findings have powerful implications for future applied work.

F-15: The Effect of Celebrity and Visual Attention on Value-Based Decisions

Simone D'Ambrogio[1], Noah Werksman[1], Michael Platt[1], Elizabeth Johnson[1]
[1]University of Pennsylvania

Marketers have long used celebrities in advertisements to help consumers build strong associations between brand and product. Celebrity endorsements appear to increase brand awareness and recall, but precisely how celebrity endorsements shape the value-based decisions remains unclear. Another factor that shapes attention in social contexts is gaze following, in which a viewer aligns their attention with the gaze direction of another individual. Thus, celebrity and gaze cues may independently and possibly jointly contribute to consumer preferences induced through advertisements. Here we investigated the effect of celebrity and gaze cues in advertisements on both gaze behavior and binary choices between competing products. We used the drift diffusion model (DDM) and the attentional drift diffusion model (aDDM) to quantify the decision process and estimated subjective value based on a set of binary decisions. We recruited seventy three adults (52 women, 21 men; mean age = 26) for this study. Subjects' eye movements were recorded using a SMI Red250 eye tracking bar. Every subject performed three tasks in the following order: 1) a two-alternative forced choice between paired combinations of 15 common snack foods; 2) passive viewing of 15 mock advertisements. Of the 15 images, 3 were advertisements of snack foods alone, 3 featured a modern celebrity gazing at the snack, 3 with a celebrity gazing at the viewer, 3 with a non-celebrity gazing at the snack and 3 with a non-celebrity gazing at the viewer; 3) a repeat of the first task. Subjects also filled a brief questionnaire about celebrity recognition and demographics. Eye tracking analyses indicate that the presence of a celebrity does not significantly increase overall dwell time on the promoted snack or face, although celebrity-endorsed ads drive significantly more visits to the face. Both celebrity and non-celebrity

endorsers gazing at the viewer increased dwell time on the face and visits toward the featured snack. Participants were more likely to choose a snack when it was promoted by a celebrity, and that effect was explained by changes in the drift rate of the DDM. In addition, we developed a new version of the aDDM which revealed that the subjective value of the snacks that were advertised with a celebrity is less influenced by visual attention than the value of snacks advertised with a non-celebrity, or snacks presented alone. These findings further our understanding of the psychological mechanisms underlying the impact of celebrities in static advertisements. Furthermore, our results shed new light on the factors that affect the impact of visual attention on value-based decisions.

F-16: Differential neural activity associated with performance monitoring in a financial trading task

Alison Harris[1], Nicolas Maramica[1]
[1]Claremont McKenna College

Objective: The ability to monitor one's performance and adaptively change behavior in response to environmental feedback is critical to complex decision-making. Yet evidence from financial trading suggests that sensitivity to past events can bias behavior away from optimality: for example, leading to a tendency to sell assets that have gone up in value relative to purchase price, known as the disposition effect. However, the link between neural measures of performance monitoring and the disposition effect across individuals is relatively unexplored. Here we used event-related potentials (ERP) to examine a midline frontal negative deflection, previously associated with adaptive control of behavior, during a financial trading task. Methods: ERPs were recorded while participants (N = 60) traded in an experimental stock market to obtain real monetary outcomes. On each trial, subjects saw an update screen with the original cost and current price of one of three stocks, and then decided whether to buy or sell that or another stock. The disposition effect was defined as the proportion of gains realized minus the proportion of losses realized across the course of the experiment. Sensors showing a significant midline frontal negative deflection following the update screen were identified from an initial subtraction of Loss - Gain across all participants, then entered into a mixed-design ANOVA with asset ownership (Held/Not Held) and price change (Gain/Loss) as within-subject factors. Disposition effect group (High/Low) was computed via a median-split analysis, and included as a between-subjects factor in the analysis. Results: Preliminary analysis identified a time window with negative midline frontal deflections from 150-250 ms post-stimulus. We found significant main effects of asset ownership ($F(1,58) = 9.95$, $p = 0.003$) and price change ($F(1,58) = 6.612$, $p = 0.02$), reflecting more negative responses for stocks that were held and for losses relative to gains. Critically, the interaction of asset ownership with disposition effect group was significant ($F(1,58) = 8.51$, $p = 0.005$): whereas the high disposition effect group responded similarly to price changes of owned and unowned stocks, the low disposition group showed a reduced frontal negativity toward unowned stocks. Conclusions: These data suggest a

connection between performance monitoring, as indexed by the frontal negativity, and decision-making in a financial trading scenario. Individuals with higher disposition effects showed similar neural responses to price changes of owned and unowned stocks, supporting a link between sensitivity to counterfactual outcomes and suboptimal choice behavior.

M-1: Cognitive sequential dependencies in the wild: A Sentiment analysis approach

Mohsen Dolatabadi[1], Javad Salehi Fadardi[1], MOHSEN KAHANI, HOSSEIN KARSHKI[1]
[1]Ferdowsi University of Mashhad

Social networks data as naturally occurring data, are nowadays freely available to researchers and can constitute a unique source for exploring many theories in psychology and cognitive sciences. This paper investigates an effect known as cognitive sequential dependence in decision making, using five-point rating in term of review polarity index (RPI) of each user's reviews by natural language processing about services or businesses in YELP. In the present study, the criteria for cognitive dependency in decisions is the degree of deviation of the present RPI from the average of each user's RPIs. The statistical population consists of all user's reviews published by the YELP site, which contains over 6 million reviews. After some initial preprocessing and filtering on textual reviews by Stanford CoreNLP tools in java, the linear regression analysis was performed on the reviews. Regression coefficients between deviation from the mean of RPIs, with the RPI, corresponded to different distances from current review up to 7 distances against a baseline, represented statistically significant and strong relationships, as well as they, revealed by going farther from current user decision, a subtle change from contrast effect to assimilation effect in users' decisions. While there are not usually subtle matches between the polarity of explicit ratings (stars) and provided textual data in business websites, The promising results of this study suggest utilizing textual content of reviews as implicit ratings to track cognitive sequential dependencies in the wild data and application of it for debasing algorithm and designing efficient online recommendation systems.

M-2: Integrating gaze predictions when fitting dynamic models of attention and decision making

Sebastian Gluth[1]
[1]Department of Psychology, University of Basel

Over the past decade, the role of attention in value-based decision making has become a central research topic in neuroeconomics. On the basis of eye-tracking data, extensions of extant sequential sampling models of decision making have been proposed, including the attentional Drift Diffusion Model (aDDM). These models take the influence of attention on preference formation into account and can thus be tested not only on choices and response times (RT) but also on eye-tracking data. Importantly, however, parameter estimation of these

models has relied exclusively on choices and RT so far. This limitation is particularly problematic in light of mounting evidence that attention itself is not allocated randomly but changes dynamically as the upcoming decision emerges. The goal of the present work is to overcome this methodological limitation and to include eye-tracking data when estimating parameters of sequential sampling models. Using the aDDM and a recently proposed extension of it, a general approach to integrative modeling of choices, RT, and eye movements is presented. This approach combines extensive model simulations with probability density approximation and differential evolution Markov Chain Monte Carlo sampling to enable hierarchical Bayesian parameter estimation. With respect to eye-tracking data, the approach focusses on fixations and takes their latencies and locations into account. The proposed integrative modeling framework is shown to provide a more sensitive comparison of different implementations of the aDDM, and it promises to substantially advance the research field on computational modeling of attention and decision making in general.

M-3: Attention and choice: A multi-stage process?

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

Westbrook et al. (2020, Science) argue that choice involves two stages, an early stage where gaze amplifies information, and a late stage where gaze is directed towards the implicitly chosen option. We argue that their eye-tracking analyses are not sufficient to support this claim because they would identify two stages even in single-stage data. Prior work has shown that attention imparts an additive bonus for the attended option, and also amplifies (multiplies) its value. In the attentional drift diffusion model, the multiplicative effect is captured by a discount on the unattended option. Westbrook et al. makes the claim that multiplicative effects occur early in the trial while additive effects occur late. One analysis shows a gaze cascade effect, where the probability of attending to the chosen option steadily increases in the time prior to the response. Another analysis examines the bifurcation of gaze, splitting each trial into early and late stages, then examining the additive and multiplicative effects within each stage. Multiplicative effects were strongest pre-bifurcation, while additive effects were strongest post-bifurcation. We provide new analyses questioning the validity of these points. We find that the gaze cascade effect occurs even in models where gaze is random. It arises from a correlation between gaze and preference, coupled with autocorrelation in gaze location. Thus, the presence of a gaze cascade is neither evidence that gaze is being drawn to the to-be-chosen option, nor evidence for additive effects. Additionally, the gaze cascade effect guarantees a post-bifurcation gaze bias. This may affect estimates of how gaze affects choice, indicating different effects pre- vs. post-bifurcation, without actual differences in the choice process. We tested this by simulating a dataset with a purely multiplicative model. We then examined pre- and post-bifurcation model fits, using different bifurcation points. We found that model parameters were highly sensitive to the selection of the bifurcation point, and so cannot be

relied on to distinguish among cognitive mechanisms. We also re-evaluated Westbrook et al.'s evidence for additive gaze effects. We re-fit their additive and multiplicative models to data simulated with purely multiplicative effects. These simulations assume a constant discount rate on all unattended elements. The results of this model comparison erroneously yielded evidence for additive models, even when they were absent. In summary, we show that Westbrook et al.'s experiment highlights that attention is largely multiplicative, and results to the contrary were due to flawed analyses.

M-4: Single-trial estimates of sequential sampling models parameters are not just noisy but can also be biased

Laura Fontanesi[1], Gilles de Hollander[2], Sebastian Gluth[1], Jörg Rieskamp[1]
[1]University of Basel, [2]University of Zurich

Traditionally, researchers estimate parameters of sequential sampling models (SSMs) from repeated choices across different conditions. Crucially, differences in parameters across conditions are interpreted as shifts in the underlying cognitive processes: For example, lower decision thresholds under high time pressure are interpreted as decreased cautiousness. Recent work has explored whether the parameters of SSMs can be estimated at a more detailed, single-trial level as well, to infer shifts in cognitive processes in subsequent trials. Such a more detailed window on decision-making processes has exciting applications. For example, by correlating single-trial estimates to neuroimaging data, we can relate specific brain areas to cognitive processes that may vary from trial to trial and not merely across conditions. The present work highlights some important limitations of such a powerful approach. First, we reproduce earlier work and show that single-trial estimates of SSM parameters are extremely noisy. We also show that single-trial SSM parameter estimates can be highly biased by the outcome of a choice. For example, single-trial estimates of the rate of evidence accumulation in incorrect choices are severely underestimated when compared to the generating single-trial parameter (and vice versa for correct choices). We will show how these problems can pollute the cognitive interpretation of single-trial parameters in joint modeling approaches (such as the regression diffusion decision model and the neural diffusion decision model) and can be exacerbated by correlations to process data. Finally, we offer a potential solution where SSMs that incorporate more information about trial-to-trial differences (e.g., stimulus or feedback properties) produce more reliable single-trial estimates.

M-5: Biophysically Realistic Decision Making Model With Binary Neurons

Marcin Penconek[1]
[1]University of Warsaw

Recent developments in neuroscience provide evidence how decision making is implemented in neural networks of the brain. In 2002, Wang developed a biophysically realistic model based on recurrent attractor network with the leaky integrate-and-fire neurons and incorporating slow NMDA-mediated synapses which are thought to play role in integrating neuronal activity. The model supports probabilistic decision making and is capable to replicate results of dot motion discrimination experiments (Shadlen and Newsome, 1996, 2001; Roitman and Shadlen, 2002). The aim of my presentation is to demonstrate an alternative model based on McCulloch-Pitts (1943) binary neurons capable to reproduce the same experimental results. Data generated by the model fits Weibull function with realistic parameters for decision accuracy as a function of stimulus coherence, in line with the experiments. The model also correctly predicts reaction time for correct and error decisions as a function of coherence. Phenomenological approach used to construct the model suggests that decision making is a network architecture phenomenon depended on the existence of attractor sets rather specific activation dynamics described by the leaky integrate-and-fire neurons. Apart from using binary neurons, the model reduces the network to excitatory neurons only. The underlying assumption is that the role of inhibitory neurons relies on providing the global feedback inhibition mechanism. The mechanism that is introduced in the model ensures persistent chaotic behavior of the network, and as a result, creates stochastic inputs into the attractor pools. It also facilitates winner-take-all competition for categorical decisions.

M-6: A single neural circuit architecture that captures divisive normalization, working memory, and winner-take-all choice

Bo Shen[1], Kenway Louie[1], Paul Glimcher[1]
[1]New York University

Objective: Circuit models of decision-making utilize two motifs to capture the dynamics of decision-related neural activities. One motif dynamically captures divisively normalized value representation, including value adaptation (Louie et al., 2014; LoFaro et al., 2014; Padoa-Schioppa, 2009; Cohen & Padoa-Schioppa, 2019). The other motif captures the dynamics of a form of working memory and winner-take-all (WTA) competition (Wang, 2002; Wong & Wang, 2006). However, both motifs are restricted in their explanatory power. For example, results from fixed-duration tasks exhibit a transition from representation to competition which appears to be under the top-down guidance of action signals (Roitman & Shadlen, 2002; Louie et al., 2011). It remains unknown how a single circuit model can capture all of these features. Here we examine how well our previously described hybrid model, which incorporates both value normalization and WTA competition, captures neural and behavioral data across different tasks. Methods: The hybrid model implements recurrent excitation and global inhibition drawn from existing models and incorporates a new element of local disinhibition that governs top-down control. The model was implemented as a dynamical system capturing firing rate responses to bottom-up signals of stimuli and top-down signals for action execution. We tested

this model in capturing neural dynamics and behavior in a range of empirical paradigms: (1) a classic reaction time task (Roitman & Shadlen), (2) fixed duration tasks (Roitman & Shadlen, 2002; Louie et al., 2011), and (3) a multiple alternative choice task (Churchland et al., 2008). Results: We show that the hybrid model captures the dynamic of WTA competition in the reaction time task, the phase transition in the fixed-duration tasks, and the quantitative effects under multiple inputs. In addition, the hybrid model fits well to empirical distributions of reaction time and choice accuracy from the reaction time task. Finally, the predicted effect of GABAergic manipulation distinguishes models with and without local disinhibition, providing neural and behavioral predictions for pharmacologic/optogenetic experiments. Conclusions: We find that a biologically-plausible decision circuit with recurrent excitation, global inhibition, and local disinhibition easily adapts to the neural dynamics from different paradigms. This model precisely captures empirical psychometric and chronometric data. These results suggest that the specific computations of divisive normalization and WTA competition, under the switch control of local disinhibition, is a core circuit motif in decision-making.

M-7: OpenMonkeyStudio: Pose tracking of rhesus macaques in an open enclosure for naturalistic foraging and decision making studies

Praneet Bala[1], Benjamin Hayden[1], Hyun Soo Park[1], Zimmermann Jan[1]
[1]University of Minnesota

The neuroscience of behavioral decision-making in rhesus macaques is generally studied using laboratory contexts in which motion is limited. An alternative approach is to allow for unrestrained movement, monitor body position, and infer decisions from pose information. We developed a novel markerless pose estimation system that we call OpenMonkeyStudio. Our system makes use of 62 machine vision cameras organized in two rows that surround a 9'x9'x8' cage that affords free movement. The cameras are mounted on an exoskeleton frame mounted around the perimeter of the cage. Cameras are synchronized via a digital pulse at 30 Hz and transmit their datastream to six custom built central processor systems. We leverage the synchronized multiview image streams to reconstruct body landmarks on macaque subjects in 3D. For each image, a convolutional neural network (CNN) is used to localize 13 body landmarks based on visual appearance. A main challenge of training such CNN lies in obtaining a large training set. Unlike existing large human datasets that have been annotated by crowd workers, collecting a comparable dataset for macaques is infeasible due to the requirement of expert knowledge and large intra-class variation. We therefore instead use a semi-supervised learning approach that enables us to explore the unlabeled data using multiview geometry in conjunction with a small set of professionally annotated data. The key insight is that it is possible to geometrically transfer the annotations in one image to the other view images without additional manual efforts (cross-view self-supervision). This allows us to propagate visual information across views and time. We demonstrate a strong recognition performance of the CNN trained by less than 4% of annotated data. Given the landmark recognition on

multiview images, we triangulate them to produce 3D landmarks given estimated camera parameters (orientation and translation). In addition to 3D landmark localization, we build a volumetric representation of individual subjects to estimate detailed pose using a 3D visual hull algorithm based on semantic segmentation. Together this approach allows us to reliably reconstruct macaques poses in our freely behaving environment.

M-9: SynthEco: A Digital Platform Linking Omics, Brain, and Society Data for Multiscale Mechanisms of Lifelong Physical and Mental Health, Disease, and Resilience

Miguel Cisneros-Franco[1], Catherine Paquet[2], Alain Dagher[1], Daiva Nielsen[1], Spencer Moore[3], Patricia Silveira[1], Shawn Brown[4], Laurette Dubé[1]
[1]McGill University, [2]Université Laval, [3]University of South Carolina, [4]Carnegie Mellon University; University of Pittsburgh

Chronic physical and mental disorders such as type 2 diabetes and depression arise from the cumulative effect of genetic predispositions and contextual exposures, as well as the individual responses to their interaction over time. An interdisciplinary, digitally-powered person-in-systems approach is thus needed to better understand these complex gene-by-environment (G*E) interactions and support better targeted intervention for both promoting health and preventing disease at different points along the lifespan. Towards this goal, we developed SynthEco, a digital platform to generate synthetic ecosystems (SE)--datasets that accurately represent a population's demographic characteristics and built environment. We first present a set of open-source tools that use iterative proportional fitting of census data to generate synthetic individuals and households that are placed in a realistic environment including school and workplace assignments, food sources, and community and healthcare organizations. These base SEs may then be enriched using datasets from population-based cohorts (e.g., UK Biobank), including whole genome sequencing, social network, neurobehavioral, and brain imaging data from individuals at various stages of the life course. However, as these SEs are synthesized from aggregated data, they provide a unique environment to share such data while protecting participants' privacy. By linking various siloed data collection efforts, SynthEco will facilitate the study of G*E interactions and neurobehavior-related variables in a way that is not typically considered in cohort studies, with type 2 diabetes, depression, and self-rated wellbeing (an important predictor of future morbidity and mortality) as use cases. Specifically, we discuss the use of polygenic risk scores (PRS) and expression-based-PRS (ePRS) to investigate the relationship between genetics, environment, health outcomes, and individual differences in the neural circuits underpinning learning, attention, and cognitive control--processes found to be altered in the context of chronic physical and mental disorders. Finally, we discuss the use of environmental measures known to influence behavioral and health outcomes, such as social networks, social capital, food environment, and indices of multiple deprivation. SynthEco will serve as a starting point for agent-based simulations to model G*E interactions across the lifespan. Longitudinal data from population-based cohorts can be used

to test and improve the accuracy of these models, which in turn can be used to simulate the impact of different interventions and support decision making in urban and health interventions planning.

M-10: Fixation patterns in simple choice are consistent with optimal use of cognitive resources

Frederick Callaway[1], Antonio Rangel[2], Thomas Griffiths[1]
[1]Princeton University, [2]California Institute of Technology

Even when choosing among a small set of alternatives, people don't perfectly evaluate every option. Instead, decisions seem to be made on the basis of sequentially accumulated, noisy evidence. Previous research has suggested that the accumulation process is modulated by visual attention such that evidence accumulates more rapidly for attended items. But what guides attention itself? To answer this question, we formalize decision making as a Bayesian evidence accumulation process---similar to a drift diffusion model, but with explicit representations of uncertainty. To capture attention, we assume that samples can only be collected for the fixated item. We additionally assume a cost for each sample as well as a cost for switching between items (making saccades). The problem of attention allocation is thus cast as a sequential decision problem in which a decision maker must continuously decide whether to select an item or keep sampling, and in the latter case, which item to sample from. The optimal attention allocation policy is the one that maximizes the expected value of the chosen item less the costs incurred by the decision making process. We approximate the optimal fixation policy using tools from metareasoning in artificial intelligence. We find that fixations are drawn to items whose value estimates are uncertain and close to those of the competing item(s). Furthermore, we find that in the case of trinary---but not binary---choice, attention is preferentially directed to items with higher estimated value. The model thus provides a normative foundation for recently proposed models of both uncertainty-directed and value-directed attention. It additionally specifies a near-optimal tradeoff between these two factors, as well as a near-optimal stopping rule and fixation-termination rule. Comparing model predictions to human behavior in two previously collected binary and trinary choice datasets, we find that the model accounts for many previously identified effects and we also confirm several novel predictions. Together, our results suggest that the evolving state of the decision process influences the fixation process in a way that is consistent with the optimal use of limited cognitive

M-11: Multi-slice FMRI acquisition compromises detection of subcortical reward responses

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

Researchers have raised concerns about the reproducibility and replicability of Functional Magnetic Resonance Imaging (FMRI) findings but have not investigated the impact of concurrent changes in data acquisition protocols on signal quality in different brain regions. Recent innovations in neuroimaging techniques have enabled the simultaneous acquisition of multiple slices of data ("multi-slice" acquisition), which has the benefit of facilitating acquisition of more data in the same amount of time, but at a potential cost of inducing artifacts (Risk et al., 2018). In a combined meta-analysis and direct comparison of FMRI findings using the Monetary Incentive Delay (MID) task, we therefore explored whether the use of multi-slice versus single-slice pulse sequences might compromise detection of reward-related activity in predicted subcortical brain regions. Meta-analytic findings suggested that relative to single-slice acquisition, multi-slice acquisition significantly decreased predicted effect sizes for gain anticipation in the Nucleus Accumbens (NAcc) as well as for gain outcomes in the Medial PreFrontal Cortex (MPFC). Further, direct analytic comparison of MID task data acquired with multi-slice (multiband factor = 6) versus single-slice pulse sequences on the same scanner (n=12) suggested that high frequency artifact was the primary correlate of signal loss in subcortical regions. Together, these findings suggest that when assessing subcortical activity with FMRI, researchers should adopt conventional single-slice rather than multi-slice acquisition protocols. These results also suggest caution in interpreting subcortical responses (or the absence thereof) in existing task-related and resting state experiments utilizing multi-slice acquisition protocols.

M-12: Costly Multidimensional Information

Ambuj Dewan[1]
[1]NYU Grossman School of Medicine

BACKGROUND & OBJECTIVE Many objects of economic interest have multiple attributes (e.g. computer specifications, job candidate qualifications) whose values a decision-maker (DM) must learn about before coming to a decision about them. Learning about those attributes can involve the use of scarce attentional resources. Using a rational inattention approach, a Bayesian efficient coding framework where the entropy reduction from prior beliefs to posterior beliefs is constrained, this project analyzes the optimal attentional strategy for learning about multidimensional objects. THEORETICAL FINDINGS We begin by assuming jointly normally distributed dimensions and quadratic loss. Solving for optimal attentional allocation identifies two notable properties: sparsity - the neglect of insufficiently variable dimensions; and summarization - the collapsing of correlated dimensions into orthogonalized summary statistics. An optimal DM pays more attention to the more variable orthogonal components. Thus, the baseline model provides an informationally optimal rationalization for intuitive attentional heuristics. When applied to a consumer choice scenario, the baseline model can also rationalize choice overload and focusing illusions in the valuation of multiattribute goods. The model is extended to account for non-Gaussian distributions of characteristics and other

types of loss, in particular Laplacian priors and absolute loss. Though many of the insights from the baseline normal-quadratic case remain, the solution's sparsity property incorporates stochasticity: the DM randomizes which dimensions to attend to, with probability increasing in prior variability. The resulting attentional strategy provides an information-theoretic rationalization for salience illusions in the valuation of multiattribute goods. The model is also extended to account for mistaken beliefs about the distribution of attributes and restrictions on which attributes the DM can directly observe. This allows the model to make predictions about the efficacy of anti-discrimination initiatives that restrict what questions an employer, landlord, contract manager, etc. may ask a potential candidate. The key result is that such initiatives are more effective when the restricted dimensions are less correlated with the observable ones, because it limits the amount of information that can be gleaned about the unobservable dimensions from the observable ones.

M-13: Normalized reinforcement learning predicts asymmetric and adaptive choice behavior

Kenway Louie[1]
[1]New York University

Novel, uncertain, or dynamic environments require organisms to learn appropriate behavior based on environmental feedback. This learning is widely modeled in psychology, neuroscience, and computer science by prediction error-guided reinforcement learning (RL) algorithms. While standard RL models assume a linear reward function, growing empirical evidence suggests that neural activity in diverse brain areas is a saturating, nonlinear function of reward; however, the computational and behavioral implications of nonlinear RL models are unknown. Here, we characterize the behavior of an RL algorithm with a specific nonlinear value function implemented via divisive normalization. We find that normalized reinforcement learning (nRL) models generate an intrinsic and tunable asymmetry in reward prediction errors (RPE), with a varying bias for negative (worse than expected) versus positive (better than expected) outcomes. Our simulations show that NRL RPE asymmetry explains diverse behavioral and neural findings previously attributed to different mechanisms, including: (1) variable risk preferences in choice under uncertainty (Niv et al. 2012, Lefevbre et al., 2017), (2) parallel dopaminergic information channels, proposed to be a neural mechanism for distributional RL (Dabney et al. 2020), and (3) adaptation phenomena in dopaminergic activity and probabilistic choice behavior (Tobler et al. 2005, Rigoli et al. 2016). These results suggest that nonlinear RL reflects an adaptive response to ecological constraints, and argue for an incorporation of biologically valid value functions in computational models of learning and decision-making.

M-14: Bridging self-report and choice modelling to investigate the accuracy of self-inferred preferences

Ryan Carlson[1], Hongbo Yu[2], Molly Crockett[1]
[1]Yale University, [2]University of California Santa Barbara

Objective: The inferences we draw about our own preferences have crucial implications for not only how we see ourselves, but also how we interact with the social world. Influential theories in social psychology suggest that our self-inferences about our internal states are largely biased and inaccurate. Little work, however, has directly tested how closely self-inferences of preferences track with actual preferences as expressed in behavior. Here, we use computational modeling to assess the accuracy of self-inferences, and whether accuracy differs for social and non-social preferences. Methods: Participants (N=71) made a series of 52 binary choices ('Option A' versus 'Option B') across two conditions. In the social condition, they chose between earning money for themselves versus The Red Cross charity. In the non-social condition, they chose between earning gift cards for Starbucks versus Home Depot. Each trial contained varying amounts of money associated with Option A versus B, ranging from 6A:1B to 1A:6B. At the beginning, middle, and end of each block, participants self-reported their inferences about the relative strength of their preferences for Option A versus B, as well as their confidence in these self-inferences. Crucially, we measured self-inferences on a scale that exactly matched the possible utility weights participants could assign to Option A versus B. This allowed us to directly compare self-inferences about Option A relative to B, with actual preferences for A relative to B. To do this, we fit participants' choices to a utility model that captured the relative subjective values for Option A versus B and compared the model weights to participants' self-reports of inferred preferences for profit relative to charity. Results: Participants expressed high confidence in their inferred preferences for the options in both social and non-social choice blocks. Retrospective self-inferences closely tracked with choice-based model estimates of their preferences, both in the social domain (rho = .78, $p < .001$) and non-social domain (rho = .79, $p < .001$). A similarly strong relationship was also observed for prospective self-inferences (social: rho = .79, $p < .001$; non-social: rho = .77, $p < .001$). Conclusions: This method of comparing self-inference and model-estimates of preferences has important implications, not just for understanding peoples' broader capacity for self-knowledge, but also for discerning when people will be more or less precise in the inferences they draw about their preferences.

M-15: Examining the Neural Correlates of Value-Based Decisions for Products in a Naturalistic Shopping Environment; a Combined Mobile EEG and Eye Tracking Investigation.

Hannah Roberts[1], John Tyson-Carr[1], Alice Newton-Fenner[1], Katerina Kokmotou[1], Adam Byrne[1], Nick Fallon[1], Timo Giesbrecht[1], Andrej Stancak[1]
[1]University of Liverpool

Previous electroencephalographic (EEG) studies in humans have shown that valuation decisions can be computed within hundreds of milliseconds (Harris et al., 2011). Despite this, research

regarding neural temporal dynamics of value-based decisions at the time of choice in realistic scenarios is limited. Building on Roberts et al (2018), the aim of the current study was to examine the spatiotemporal dynamics that underpin economic decisions for products in a realistic shopping environment using mobile EEG and eye tracking techniques. Twenty-seven participants made purchasing decisions about 216 household products while mobile EEG and eye tracking was recorded. Products were divided into three retail value categories (£0.50-£12) and three products from each category were displayed on all 36 shelves. A Becker-DeGroot-Marschak auction elicited willingness to pay for every product, creating four subjective value (SV) conditions. Guided saccade topographies were clustered using a Principle Component Analysis (PCA) and subtracted from the data using a regression. Eye Movement Related Potentials (EMRPs) were examined and an Independent Component Analysis separated neural activation from grand averaged EEG. A clustering solution was applied to grand averaged EMRPs using the PCA method. Each cluster's mean activation was statistically compared across SV conditions. Four clusters were modulated by SV. Enhanced amplitude was found for low medium and high compared to low (Cl6 73-83 ms) and low and low medium compared to high value (Cl4 129-139) within the latency of the lambda component. Enhanced amplitude was found for low medium and high compared to low (Cl5 170-183 ms) and low, low medium and high compared to high medium value (Cl3 217-250 ms) within the latency of the P200 component. Finally, enhanced amplitude was observed for all SVs compared to low value in the latency of the N400 (Cl3 324-359 ms). To our knowledge, this is the first study to successfully examine EMRPs underlying value-based decisions for products in a naturalistic shopping environment. In accordance with Roberts et al. (2018), low value products appear to receive neural prioritization in real world scenarios as they are isolated and contrasted across all clusters. High value products receive initial prioritization of attention and are contrasted against lower value items (within the lambda latency). Later, medium value products are isolated and incrementally compared as they require more deliberation. Early prioritization of attention towards high value products and then low value products could represent a binary approach-avoidance mechanism that exists to maximize economic resources.

M-16: Modifying the Trust Game: a meta-analysis

Zuzanna Tajchman[1], Shulang Yue[1], Iris Vilares[1]
[1]University of Minnesota

Over the last 20 years, the Trust Game has become the go-to paradigm to study trust, and has been employed in experimental studies in psychology, economy, neuroscience and more. Its common application gives us opportunity to examine to what extent altered experimental protocols, although build on a similar structure, have an impact on the results of the trust game. In our metanalysis, we decided to investigate how some of the most commonly modifiable factors influence the outcomes of the Trust Game. One of the objectives of the meta-analysis was to analyse how the number of trials played, with and without the same

partner throughout the experiment, effects the trust game results. Additionally, we examined how participants played when they believed that the other player was human versus when they knew they were interacting with a computer only. Another factor that can determine behavioral strategies in the TG is the "rate of return"- amount by which the money is multiplied from player one to player two. Finally, we looked at how different degrees of interaction with the other partner and how and how interdependence between compensation and performance affected behavior. Literature search for this meta-analysis was conducted by two independent researchers in three databases: PubMed, Web of Science and Embase. We used as search terms "trust game" or "investment game" as search terms. The final selection included 47 studies used in the meta-analysis. Within those 47 included articles, 42 provided results for the amount sent ("trust) and 22 observation for the amount returned ("reciprocity). Our results showed a significant negative relationship between number of trials and trust, when playing with different partners. Results also showed that participants trusted more when they believed they were playing with a computer vs if they believed they were playing with a human counterpart. In addition, participants trusted less if given a flat rate for their participation in the study vs. if payment was based on performance/behavior in the game. The other factors studied showed no significant effects, and we did not find significant effects on reciprocity. Our results suggested that modifying the TG protocol could affect trust. Number of trials plays a crucial role in building general cooperation through the game, but surprisingly this effect was negative. Moreover, it appears that participants are more likely to cooperate when they are not playing with (or believe to be playing with) a human counterpart. Further, our meta-analysis supported previous evidence that when participants' payoff depends on behaviour in the game, their trust is higher.

M-17: Capturing Process-Level Predictions of Diffusion Models with Mouse-Tracking

Nitisha Desai[1], Paul Stillman[2], Kentaro Fujita[1], Ian Krajbich[1]
[1]The Ohio State University, [2]Yale School of Management

When people make decisions, they may display initial biases, then waver between the options, as assumed by the drift diffusion model (DDM). The DDM treats these as hidden processes, estimated from choice and RT distributions. But are these real processes? If so, can they be measured directly? We show that DDM parameters do correspond to features of subjects' mouse movements. We explored this relationship in an intertemporal choice (Stillman, Madvedev, & Ferguson, 2017) and a risky choice (Stillman, Krajbich, & Ferguson, in press) dataset. In the first experiment, subjects (n=140) completed 180 decisions, choosing between a smaller amount of money today (SS) and a larger amount later (LL). In the second, subjects (n=253) made 215 decisions, choosing between a 50/50 gamble and a certain option. We hypothesized a negative correlation between the trial-level drift magnitude and area under the curve (AUC), i.e. the directness of the trajectory. We also predicted that starting-point bias would relate to the initial angle of the average trajectory. To test these predictions, we fit

DDMs to our choice and RT data using the software HDDM (Wiecki, Sofer, & Frank, 2013). We found a significant relationship between AUC and drift rate in both datasets: the smaller the AUC (indicating a straighter path), the larger the drift rate (intertemporal: p=10-16; risky: p=10-16). These results are expected, as a larger drift rate signifies a stronger preference for one option and thus less conflict between them. In both datasets, we also found a correlation between starting point and initial angle of the trajectory (intertemporal: r(125)=0.596, p=10-13; risky: r(186)=0.237, p=0.001). For example, subjects with a starting point bias towards SS had an initial angle biased towards SS. Lastly, we used mouse-tracking to understand fast vs. slow errors. DDMs predict slower errors relative to correct responses when the starting point is biased towards the correct response and faster when biased towards the error. Using the initial angle of trajectory as a proxy for starting point, we predicted errors (i.e., choosing the lower utility option) to be slow if the angle pointed towards the correct option, but fast otherwise. Thus, we investigated how the correct-error RT difference covaried with direction of the initial angle. In both datasets, we found a significant interaction between direction of the initial angle and response (correct vs. error) on RT after controlling for utility (intertemporal: β=-0.238, p=9.42×10-6; risky: β=-0.129, p=1.30×10-7). Our work shows mouse-tracking elucidates process-level predictions of the DDM.

N-1: Oxytocin improves delay of gratification and cognitive flexibility in non-social decision making

Georgia Eleni Kapetaniou[1], Matthias Reinhard[1], Patricia Christian[1], Andrea Jobst[1], Philippe Tobler[2], Frank Padberg[1], Alexander Soutschek[1]
[1]Ludwig Maximilian University Munich, [2]University of Zurich

Oxytocin has been widely studied for its effects on social cognition and decision making. Recent findings, however, challenge the specificity of oxytocin for the social domain and suggest a domain-general allostatic function for oxytocin through promoting future-oriented and flexible behavior. Here, we tested the impact of intranasal oxytocin on these two core aspects of non-social decision making, delay of gratification (as measure of future-oriented behavior) and cognitive flexibility, i.e. the capacity to re-learn dysfunctional associations between environmental cues and outcomes. We conducted a pre-registered randomized, placebo-controlled, within-subject study in which 49 healthy male participants performed a battery of tasks measuring delay of gratification (intertemporal choice task), cognitive flexibility (reversal learning) and inequity aversion (dictator game) after administration of either oxytocin (24 IU) or placebo. The participants also performed a working memory task both before and after substance administration. In the intertemporal choice task, model-free and model-based analyses provided converging evidence for reduced impulsivity under oxytocin relative to placebo. Oxytocin further improved reversal learning, however only for individuals with low baseline working memory performance. In the dictator game, oxytocin increased generosity only under conditions of advantageous but not disadvantageous inequity, extending previous

findings for the role of oxytocin in social decision making. Finally, we found no evidence for oxytocin on working memory performance, speaking against a mediating role of working memory for the observed oxytocin effects on decision making. Taken together, our findings show that oxytocin affects both social and non-social decision making, questioning the specificity of oxytocin for the social domain and supporting recent theoretical accounts of domain-general functions of oxytocin. This may have clinical implications for the use of oxytocin as pharmacological agent in psychiatric disorders.

N-2: Repugnant Warnings, Addiction, and Rational Choice

Mariana Gerstenblüth[1], Jeffrey Harris[2], Patricia Triunfo[1]
[1]Social Sciences Faculty,. Universidad de la República, [2]MIT

It has been frequently observed that aversive stimuli and negative emotions interfere with rational decision-making (Leith and Baumeister 1996, Luce 1998, Bechara et al. 1999, Lerner, Small, and Loewenstein 2004, De Martino et al. 2006, Hewig et al. 2011, Guclu et al. 2012). We design here an experiment to specifically test whether smokers' risk perceptions are consistent with these general observations. We ask current smokers to repeatedly choose which of two experimental cigarette packages with varying warnings and background colors is less risky. We supplement our binary choice data with measurements of subjects' eye movements and time to response. Confronted with repugnant, threatening images - which included a dead fetus, a cadaver, an ulcerated tumor - participating smokers nonetheless made choices that were context independent, adhered to transitivity, and consistent with an additive utility model. Eye tracking measurements confirmed that the choices of 65 percent of participants were further compatible with a noise-reducing lexicographic utility model. This subset of participants smoked significantly more cigarettes per day. Our findings support a model in which addiction permits the smoker to suppress aversive stimuli and negative emotions that would otherwise interfere with short-term rational decision making.

N-3: Effects of memory for emotional valence and episodic detail on decision-making in aging

Michael Cohen[1], Karolina Lempert[1], Kameron MacNear[1], Frances Reckers[1], Laura Zaneski[1], David Wolk[1], Joseph Kable[1]
[1]University of Pennsylvania

Objective: Older adults often show an increased vulnerability to fraud and frequently have decline in episodic memory. We examined how memory for a single prior experience influences older adults' decision-making. Methods: We presented older (n = 41) and middle-aged (n = 31) adult participants with pictures of social partners in a "dictator game", or of houses in an analogous non-social task. Participants could gain $5 (high-reward) or $0 (low-reward) from each face or house. Each image was shown once. Minutes later, participants were asked to

choose whether to play again with each person or house. After choices, memory for each face/house and for associated reward was tested. Neuroimaging data were collected from a subset (n = 11) of older adults. In a behavioral follow-up study (n = 137), collected online and sampled across the lifespan, we examined whether choices are driven by memory for reward valence (i.e., high vs. low reward) or specific episodic details. Here, rewards were varied in 10-cent increments within each valence (low: $0-$1.50; high: $3.60-$5.00), and the memory test included the exact value (e.g., $3.80), a same-valence incorrect lure (e.g., $4.30), and 2 opposite-valence values (e.g., $1.00). Results: While older adults showed poorer associative memory overall, accuracy of confident responses was comparable across age groups. Confident memory for face/house-reward ($0 or $5) associations was related to subsequent choices in both age groups, as participants were much more likely to play again with faces or houses that they correctly believed were high vs. low value, though the effect was somewhat weaker in older adults. A preliminary analysis of fMRI data suggests that making adaptive choices (playing with rewarding partners or avoiding non-rewarding partners) was associated with increased BOLD signal in amygdala and hippocampus/parahippocampal cortex, relative to maladaptive choices. In the behavioral follow-up study, choices were consistent with memory responses even when the same-valence but incorrect lure was chosen. Memory for exact values was still above chance, however, and adaptive choices were even more likely when exact memory was available. Conclusions: Older adults are capable of using explicit associative memory to guide decisions. However, across the lifespan, such decisions may rely largely on emotional valence memory (reliant on amygdala), which tends to be more preserved with age than episodic memory for details (e.g., Limbert et al., 2018). The latter appears to provide a further boost to decision-making, however, so age-related memory deficits could still lead to suboptimal choices.

N-4: Identified motivation predicts time spent on homework during the COVID-19 lockdown across different middle school grades.

Léa Combette[1], Jean-Yves Rotgé[1], Liane Schmidt[1]
[1]Brain and Spine Institute

On the 16th of March 2020, French schools suddenly closed because of the COVID-19 pandemic. For middle schools, around three million students were then asked to study from home with no real interactions with teachers or classmates for weeks. In this very unusual context, we aimed to identify the psychological factors that underline inter-individual differences in engagement in schoolwork assignments. Research from social psychology has shown that implicit theories of intelligence (whether one believes that intelligence is a malleable state or a fixed trait), and type of motivation (intrinsic, extrinsic, or identified) influence learning and student well-being. However, it is unknown how these factors contribute to (1) inter-individual differences in time spent on homework without the physical presence of teaching staff, and (2) in the context of experiencing an adverse life event with uncertain

outcome such as the Sars Cov2 pandemic. To test these questions, we conducted an online survey study across 98 middle school students (age =13 +/- 1.2 years) from three different middle schools in the Paris area. A linear mixed effects regression showed that identified motivation positively predicted the time spent on math assignments (ß =1.04, SE =0.36, p = .004) with no significant effect of extrinsic or intrinsic motivation, implicit theory of intelligence and optimism. The more students thought that doing homework during the lockdown was important and useful for their future life, the more time they spent time studying. We also found a negative effect of grade (ß =-0.77, SE =0.31, p = .01) indicating that students in earlier grades spent more time on homework than older students. Importantly, an out-of-sample cross validation found that identified motivation in early grade students (6th and 7th grade) significantly predicted how much time the 8th and 9th grade students spent on homework (Pearson's R observed over predicted study time in N= 51 8rd and 9th grade students: = 0.3, p = .03, 95% CI due to chance: -0.24-0.22). These findings suggest that identified motivation was a general determinant of school engagement across different grades and students in middle school. We will continue data collection at the beginning of the new school year in September 2020 to test how specific these findings are to the COVID-19 lock down context.

N-5: Validating the Roles of EEG, Heart Rate, and Galvanic Skin Response in Measuring Emotions, Memory and Ad Effectiveness

Davide Baldo[1], Rich Timpone[1], Vinod Venkatraman[2]
[1]Ipsos, [2]Temple University

OBJECTIVES: Despite the growing popularity of neuroscience-based measurements in marketing, their effectiveness and complementarity remain unclear. Here, we validated the specific roles of electroencephalography (EEG), heart rate (HR) and galvanic skin response (GSR) in measuring emotions and memory. We also sought to extend their implications for predicting advertising effectiveness. METHODS: A total of 93 participants (42 female, mean age = 33.2, SD = 7.6) completed this lab study in NYC. We recorded EEG (Biosemi, 32 channels), GSR, and HR data in addition to self-report measures as respondents were exposed to an in-house emotional dataset consisting of 20 TV ads, 49 images and 16 videos across separate blocks. These stimuli were validated apriori to span a broad range of emotional valence and arousal ratings. Stimuli were randomized within each block. Participants also completed 26 trials of a gambling task aimed at inducing different levels of valence (win, lose) and arousal ($1 vs. $3), followed by a surprise recognition task for TV ads. Finally, we also obtained an independent out of sample ad effectiveness measure for the TV ads. This ad effectiveness measure is a composite score that incorporates purchase intent and brand recognition components, and has been validated to reflect in-market performance. RESULTS: The interbeat interval (IBI) consistently predicted emotional valence: IBI was larger (slower heart beat) for negative than positive images, negative than positive videos, as well as losses than gains. The frontal alpha asymmetry (FAA) from EEG was also a predictor of valence in the gambling task, but not for the image stimuli.

Finally, the skin conductance amplitude was a significant predictor of arousal across images and videos. In terms of memory, only IBI during encoding of TV ads predicted recognition accuracy. Critically, while self-reported valence and arousal did not predict out of sample ad effectiveness, both FAA and IBI were significant predictors. Moreover, FAA and IBI jointly lead to a significantly better prediction of ad effectiveness than each measure independently, indicating their complementary value. Specifically, we found that a double dissociation with FAA being linked to the purchase intent while IBI was related to brand recognition component of the ad effectiveness. CONCLUSIONS: Our findings demonstrate distinct but complementary roles for the different methods in measuring emotions and memory. Critically, IBI and FAA play a complementary role in predicting ad effectiveness. These findings highlight the value of multichannel studies and provide valuable insights into the dynamics of decision making across domains.

N-6: Heritability of belief updating is valence-dependent

Valentina Vellani[1], Neil Garrett[2], Anne Gaule[1], Kaustubh Patil[3], Tali Sharot[1]
[1]UCL, [2]Oxford University, [3]Jülich Research Center, Heinrich-Heine-Universität

Different rules and mechanisms underlie updating of beliefs in response to positive and negative information. Here, we examine whether the heritability of belief updating is also valence dependent. To that end, we tested monozygotic (MZ) and dizygotic (DZ) twin pairs on a belief update task (N = 528). MZ twins are identical twins whom share 100% of their genes. DZ twins are non-identical twins whom share 50% of their genes. Heritability can be assessed by comparing how similar MZ and DZ twins are in relation to a phenotype. We used this approach to estimate the heritability of belief updating in response to negative and positive information. The comparison of the two then allowed us to test whether heritability of belief updating is valence dependent. Our results indicate that heritability of belief updating is valence dependent. Genetic differences played a greater role in accounting for individual differences in belief updating in response to negative information (estimated heritability =11%) than positive information (estimated heritability = -6%) (difference between the two p = 0.029). In fact, no genetic basis was found for the latter. The identification of a genetic contribution to valence-dependent belief updating sits nicely with previous findings showing that individual differences in this domain are related to variations in neuroanatomy and function along the frontal-striatal pathways, which are thought to be partially heritable. By far the largest amount of variation, however, was explained by unshared environment (belief updating in response to positive information = 82%, negative information = 87%). The results provide novel evidence for a valence-dependent dissociation of belief updating and highlights the importance of both context and genetics in belief formation.

N-7: Assessing COVID-related Stress: Mental Health and Decision-Making

Candace Raio[1], Alexandra Mellis[1], Angelica Vina-Albarracin[1], Paul Glimcher[1]
[1]New York University Grossman School of Medicine

OBJECTIVE: The COVID-19 pandemic is an unprecedented and pervasive stressor that poses a direct threat to individuals' health and survival, and can also have downstream effects on mental health and decision-making. Understanding the effects of psychological stress, financial hardship and COVID-related health challenges will be critical for informing interventions throughout the pandemic. Further, gaining a more precise understanding of which factors protect against COVID-related stressors will be critical to fostering stress resilience throughout the pandemic. METHOD: To comprehensively assess these effects, we are currently conducting a large-scale, mobile (smartphone) app-based, longitudinal study to track how COVID-related stressors affect psychological health (stress levels, anxiety, depression), economic health, COVID status, and economic decision-making (risk and ambiguity preferences, impulsivity). Data collection began in early April primarily across NY State. Each day, participants spend 5-10 minutes completing gamified tasks and surveys on their smartphones, which allows us to track dynamic changes in the emotional state of participants as they navigate the challenges imposed by the pandemic. To measure COVID-related stress, we developed the COVID Stress Assessment, a weekly self-report survey that identifies general and domain-specific stressors (e.g. familial, professional, financial, physical/mental health, etc.) as well as how unpredictable, uncontrollable and overwhelming each are perceived to be. RESULTS: We examined initial trajectories of subjective stress responses and how they relate to mental/physical health and decision-making. Overall, higher perceived controllability, predictability and coping were all negatively correlated with global stress levels, suggesting they serve as protective factors that buffer perceived stress. In terms of physical health, these protective factors were negatively related to the severity of COVID symptoms participants had experienced. We further observed that higher perceived control, predictability and coping were all associated with decreased depression and anxiety as measured by the PHQ-8 and STAI, respectively. While no effects of these factors were observed on risk and ambiguity preferences, we did find that stress resilience was related to lower impulsive choice as measured by our delay discounting task. CONCLUSION: Our results provide an initial assessment of subjective stress responses, as well as resilience factors, during the COVID-19 pandemic. Our findings set the stage for future analyses tracking how emotional state predicts changes in health and choice behavior across the pandemic.

N-8: Medial frontal cortex activity predicts information sampling in economic choice

Paula Kaanders[1], Hamed Nili[1], Jill O'Reilly[1], Laurence Hunt[1]
[1]University of Oxford

Objective Decision-making not only requires agents to decide what to choose, but also how much information to sample before committing to a choice. The neural signature of voluntary

information search in economic choice has rarely been studied. Mediofrontal activity has been associated with many aspects of decision-making, such as conflict, task difficulty and belief confirmation. In this work, we are suggesting that the reason this region is encoding so many different task parameters, might be that they can all be used to guide future information search and that mediofrontal cortex is driving this behavior. Methods 30 human participants performed a decision-making task inside an MRI scanner, in which they repeatedly had to choose between two options, each consisting of two cues: a reward probability cue and a reward magnitude cue. All four cues were initially hidden at the start of a trial and two of them were sequentially revealed to the participant. After that, participants could choose to spend some of their points to view the other cues, or to make a choice between the options straight away. We analyzed brain activity in response to presentation of the first two cues. Results In response to presentation of the second cue, we found an effect of task difficulty on activity in mediofrontal cortex in a mass univariate analysis. Crucially, when we then included a predictor in this model describing whether or not participants would sample more information before committing to a choice, this main effect of task difficulty disappeared. We propose that mediofrontal cortex encodes task parameters such as task difficulty that are predictive of choice uncertainty in order to guide future information sampling. We believe the reason that this hasn't previously been shown is that most paradigms studying mediofrontal function don't allow the participant the possibility to freely sample information, making it impossible to study the role of mediofrontal cortex in information sampling. Conclusions We used an economic decision-making task together with fMRI in human participants to show that mediofrontal activity guides information sampling behavior and can account for the effect of task difficulty.

N-9: Evidence for two ways of seeing the value of multi-attribute objects

Juliette Ryan-Lortie[1], Gabriel Pelletier[1], Matthew Pilgrim[1], Lesley Fellows[1]
[1]Montreal Neurological Institute, McGill University

Many real-life decisions involve multi-attribute options. Current views are that the values associated with attributes are integrated to form a global option-value, but it is unclear whether and how the mechanisms underlying complex object recognition and evaluation come together. We recently found that vmPFC damage impaired decisions between objects when value was predicted by the unique configuration of attributes (Configural), but spared decisions when value could be assessed by combining individual attribute-values (Elemental) (Pelletier & Fellows 2019, J Neurosci). This dissociation raises the possibility that object processing is affected by reward experience, i.e. depending on whether reward is predicted by attributes or whole objects. Objective. We asked whether value assessment of multi-attribute objects in Configural and Elemental conditions engages distinct mechanisms, reflected in reaction time and eye movement differences. Methods. In two experiments, participants first learned the value of multi-attribute pseudo-objects in Configural or Elemental conditions. In Experiment 1, reaction time and eye tracking patterns were assessed while participants estimated the value of

objects presented one at a time, at different viewpoints from the learned (canonical) view. Prior research has shown that configural object recognition is slowed down by such viewpoint manipulations, whereas elemental recognition is not. This effect can serve as a behavioural indicator of the type of object representation involved. Experiment 2 asked whether fixation patterns to attributes and transitions within and between objects differed in two-option decisions under Elemental or Configural value conditions. Results. Across the two experiments, we found differences in eye movement patterns between conditions. When evaluating single objects, more transitions were made between attributes within objects, and fixations were shorter in the Configural compared to Elemental condition. Similarly, during binary choices, there were more within- than between-object transitions in the Configural condition. We further found that viewpoint manipulations increased valuation reaction time only for objects learned in the Configural condition but not those learned in the Elemental condition. Conclusion. This work provides evidence that object recognition mechanisms are differently engaged under different value assessment conditions, supporting two modes of multi-attribute object evaluation. This is in line with the known different pathways for representing elemental and configural characteristics of complex objects, and may enrich neuroeconomic models of multi-attribute decision-making.

N-10: Evidence Accumulation Under Stress

Laura Globig[1], Kristin Witte[2], Gloria Feng[3], Tali Sharot[1]
[1]UCL, [2]Max Planck Institute for Human Development, [3]Yale University

Objective: To make good decisions people often need to accumulate information under stress. This is true for personal and professional decisions, such as a patient selecting between medical treatment options, a pilot deciding how to respond to midair malfunction, or a driver stuck in traffic late for an appointment. We know little of how stress alters the process of evidence accumulation. Here, we manipulate stress levels to examine how it changes the way people accumulate information and make judgments. We hypothesized that the impact of stress on evidence accumulation is valence-dependent. In particular, as stress signals a high-risk environment, evidence accumulation may increase selectively for negative information, leading to more conservative judgements. Methods: Participants (N = 91) completed a sequential sampling task in which they were to judge whether they were in a good environment (which was associated with greater rewards) or a bad environment (which was associated with greater losses). Participants were incentivized for accuracy. Prior to the task participants in the 'stress group' (N = 45) experienced a social-threat manipulation that resulted in increased anxiety relative to participants in the 'control group' (N = 46). A classic psychometric approach and Drift-Diffusion Models (DDM) were used to tease apart the specific components of the accumulation process that were altered by stress. Results: We found that stress selectively boosted accumulation of undesirable evidence. Under stress participants required less evidence to reach the conclusion that they were in a bad environment relative to control. In contrast, the

amount of evidence required to reach the conclusion that one was in a good environment did not differ under stress and control. Drift-Diffusion Models revealed that the change was a result of a selective increase in the drift rate of negative evidence relative to positive evidence. Interestingly, while participants in the control group showed a positivity bias - they falsely believed they were in the good environment - under stress this bias was absent. Conclusions: The findings demonstrate that the process of evidence accumulation is sensitive to stress in a manner that may be adaptive. Humans flexibly adapt to environmental demands, responding to threatening environments by selectively reducing the amount of evidence required to reach undesirable conclusions. This in turn can lead to decreased likelihood of selecting risky actions. The results can help predict the impact of information on people's judgements under stressful public events such as market collapse, political turmoil and global health threats.

N-11: Computational marker of individualized learning: Predicting preference modification using Bayesian modeling of training

Tom Salomon[1], Nathaniel Daw[2], Tom Schonberg[1]
[1]Tel Aviv University, [2]Princeton University

Preference modification procedures are often characterized by significant individual differences; while some individuals demonstrate rapid behavior change, others require more time and effort. Understanding these individual differences could be beneficial for improving the efficacy of behavioral change paradigms. The Cue Approach Training (CAT) paradigm is a unique protocol that induces lasting preference modification, using a simple association of stimuli with a cue and speeded response. In the current work, we hypothesized that well-learned cue-response association in the CAT task, and ultimately stronger post-training preference change, could manifest in the form of faster anticipatory reaction times (RTs) during training. Using Bayesian modeling of CAT task's behavior, we aimed to estimate a computational marker for learning, which could potentially serve as predictive indicator for preference modification. Methods. In a meta-analysis with N=828 participants from 29 CAT experiments, we developed a Bayesian model of CAT responses. Our parameter of interest was an individualized training-dependent slope parameter, which characterized the speed and strength of an individual's shift, with training, from slow cue-dependent responses to fast anticipatory responses. We used this slope parameter as a marker for individualized learning and examined its association with individual differences in a subsequent preference modification probe task. Following the meta-analysis conclusions, we applied the model in two new studies (n=20 and n=59 preregistered replication) with a novel training regime. These studies introduced both easy and challenging CAT conditions to better exercise learning. We hypothesized that the easy training condition would be characterized by faster learning and result in more robust preference modification. Results. In the meta-analysis, across all experiments, we found that CAT induced a preference modification effect, which was associated with the individualized learning parameter. In both new studies, as expected,

individual learning parameters were more robust for the easy training condition, and preference modification effects were also stronger. Furthermore, the individualized learning marker was predictive of the subsequent preference-modification effect in both conditions. Conclusions. Our results demonstrate that Bayesian modeling of RTs in nonreinforced training with CAT can be used to identify markers of individualized learning underlying preference change. We propose that computational framework can be a powerful tool both for model fitting as well as prediction of future effect on preferences modification.

N-12: Inclusion of Neuroeconomics in Alternative and Adaptive Educational Contexts

Annie Rappeport[1]
[1]University of Maryland

Experts from neuroscience, economics and education are expressing a need to come together. Only through inclusion, participation and collaboration can new interdisciplinary knowledge be produced and used in ways that promote sustainable results for a healthier community. This need is of increasing importance given the global challenges presented by the current COVID-19 pandemic that is threatening public health, educational and economic systems. The global crisis has amplified the issues surrounding inequity and many of the most vulnerable are receiving the brunt of negative consequences. The focus of this paper are concrete ways that academia can contribute through embracing interdisciplinary collaboration to support the common good for humanity and the planet. Although this paper will focus primarily on the collaboration between economics and neuroscience, there will be inclusion of the education sector and the blossoming interdisciplinary work of neuroeducation research. More specifically, the author will connect concepts from an in-depth literature review of these fields as related to crisis response and adaptation. The researcher is heavily involved at her own university with committee working to shape behaviors of an R1 public research institution seeking to adapt during the COVID-19 pandemic. Neuroeconomic research can help people learn how their own brains work. As Trout (2009) remarks about self-awareness and biases, "Biases are also unyielding. When we are made aware of them, we don't have the ability to resolve or control them. But perhaps we can think around them--use our problem-solving powers as a means to correct them" (p 85). This idea shows realistic optimism. Neuroeconomic research has great potential at the individual level as it can assist us in being more aware of ourselves and the manipulated designs and environments around us. This awareness can allow for filtering of information and less vulnerability to external strategic misrepresentations and internal biases that may be impeding decision making. Many areas in private business are already using what is demonstrated to be predictable biases and irrational behaviors to help market and increase profits. It may be time to implement these same tools from choice architecture, nudging, and theories or cooperation to aid the mission of the public common good. Those working in neuroeconomics and neuroeducational research have the potential to support the creation of

an alternative economic system oriented to planetary and human health and can share knowledge to positively impact intentional designs during times of crisis and emergency.

N-13: Ventral striatum dopamine transporter availability is related to fMRI activation during reward receipt but not anticipation

Annie Ilsley[1], Jaime Castrellon[2], David Zald[2], Gregory Samanez-Larkin[1]
[1]Duke University, [2]Vanderbilt University

The dopamine transporter (DAT) is present on striatal presynaptic dopamine terminals and plays a role in the release and reuptake of dopamine to support reward anticipation and consumption. Few studies have tested associations between direct measures of DAT availability and reward anticipation and consumption in humans. Here, we directly measured DAT availability in a sample of 23 young adults (ages 20-29) and 23 older adults (ages 50-65) who underwent a positron emission tomography scan with the radiotracer [18F]FE-PE2I. On another visit, these participants performed the monetary incentive delay (MID) task during fMRI scanning. We first tested associations between DAT availability in the ventral striatum (VS) and fMRI activation in two predefined regions of interest (ROIs) that are critical to reward anticipation and receipt: the VS and the ventromedial prefrontal cortex (vmPFC). All analyses controlled for age and gender as covariates. Individual differences in VS DAT availability was related to VS activation during reward anticipation (t(42)=-2.290, β=-0.367, SE=0.050, 95% CI= [-0.691, -0.0435], p=0.027) and receipt (t(42)=-2.235, β=-0.361, SE=0.054, 95% CI= [-0.687, -0.0350], p=0.031), however, these effects did not survive correction for multiple comparisons across the two ROIs. We found evidence for a negative association between VS DAT availability and activation of the ventromedial prefrontal cortex (vmPFC) during reward receipt (t(42)= -3.422, β=-0.510, SE=0.033, 95% CI= [-0.811, -0.209], p=0.001), but not anticipation (t(42)=-0.683, β=-0.113, SE=0.050, 95% CI= [-0.445, 0.220]. An exploratory voxelwise analysis revealed a network of regions that were negatively correlated with VS DAT availability and reward receipt as well as a cluster in the cerebellum during anticipation. Overall, the results suggest that lower DAT availability contributes to neural processing of reward receipt but not anticipation.

N-14: Anger invigorates pursuit of immediately rewarding actions of high Pavlovian value at the expense of advantageous goal-directed actions

Alison Schreiber[1], Daryl Cameron[2], Michael Hallquist[1]
[1]University of North Carolina at Chapel Hill, [2]The Pennsylvania State University

The effects of emotions on prosocial decision making are well-described (e.g., the Empathy-Altruism Hypothesis; Batson, 1990), yet little research has specifically focused on the role of anger. Motivated by recent research on the importance of anger in promoting positive social

changes (e.g., outrage mobilizing social movements; Spring et al., 2018), we sought to examine how anger, relative to other emotions, altered prosocial decision making. Building on evidence that anger is associated with decisive action (Lerner and Tiedens, 2006), we sought to test how anger alters the balance of Pavlovian and goal-directed prosocial decision making. In this study, 27 adults completed a novel moral decision-tree task (extending a well-established deterministic decision-tree task; Huys et al., 2012) in which participants were instructed to 'help end an ongoing international conflict' by choosing actions that could either 'help' or 'hurt' the hypothetical cause. Participants were placed randomly in the action space at the beginning of each of the 48 epochs (with 4-moves per epoch; total of 192 decisions). To win points, participants needed to learn (a) the value of each action, (b) to navigate to highly valued actions in a few moves, and (c) to balance the immediate value of an action against the cumulative value of a multi-decision sequence (e.g., choosing an action with lower value because it affords opportunity to choose to higher value actions later). To assess the effect of anger, participants received an anger induction partway through the experiment. Additionally, half of all possible actions were anger-congruent (e.g., "send an angry tweet"). Mixed-effects analyses demonstrated that individuals exhibited greater difficulty learning to choose the action that could lead to the optimal multi-decision sequence if more moves remained within the epoch ($b = -1.68$, $p < .001$), suggesting that individuals "pruned the decision tree". For trials in which the optimal action was not also the immediately valuable action, we further observed that, for the 13 individuals who received the anger induction during learning (six epochs into the experiment), anger induced a tendency to choose the immediately valuable action over the optimal action ($b = -0.81$, $p = .003$), particularly if the action was anger-congruent ($b = 1.98$, $p = .006$). We further plan to conduct computational model-based analyses to verify that anger altered the balance of Pavlovian and goal-directed decision systems on choice. These results suggest that anger is likely to specifically invigorate prosocial behaviors that are immediately rewarding, particularly if the prosocial behavior is anger-congruent.

N-15: The role of the human amygdala subdivisions in prediction error encoding.

Emilia Kolada[1], Mateusz Wilk, Krzysztof Bielski[1], Izabela Szopa[2], Sylwia Adamus[1], Marek Wypych[1], Iwona Szatkowska[1]
[1]Nencki Institute of Experimental Biology PAS, [2]Neuro Device Group

Learning associations entails predicting an outcome stimulus based on a cue. When there is a mismatch between prediction and outcome, a prediction error (PE) signal is generated in the brain. Several studies have reported PE signal in the amygdala. Yet, the amygdala is not a homogeneous structure, but comprises a number of nuclei, which can be grouped approximately into two complexes: centromedial (CM) and basolateral (BL). Due to low spatial resolution of fMRI, so far only few studies have raised the issue of PE encoding in different parts of the amygdala. To contribute to the discussion on the roles of CM and BL in PE processing, we run an fMRI experiment with an adapted classical conditioning task consisting of appetitive and

aversive runs. A naturalistic reinforcer combined of gustatory (pleasant, unpleasant or neutral liquid) and visual (a video of a person drinking this liquid) stimuli followed a visual cue in a probabilistic manner. Participants had to observe the relationship between a cue and a reinforcer and anticipate the reinforcer type. Behavioral responses were passed on to the Rescorla-Wagner model, which outperformed Temporal Difference Learning and Hierarchical Gaussian Filter models, to compute PEs. To be able to precisely assign the signal to the amygdala parts we created masks of CM and BL subdivisions based on functional connectivity and Recurrence Quantification Analysis. The ROI analysis revealed that the BOLD activity in the CM subdivision correlated with PE in both appetitive and aversive learning, but only when PE was related to omission of a salient (pleasant or unpleasant) stimulus. Interestingly, no activity was found in the BL subdivision. Our results demonstrate that the CM subdivision of the amygdala is indispensable for associative learning and encoding of PE. The work supported by the National Science Centre, grant no. DEC-2014/15/B/HS6/03658.

N-16: Tracing multiscale effects of dopaminergic networks and environmental quality on later food choice and body weight: Novel expression-based striatal DRD4 polygenic risk score and evidence of differential susceptibility.

Kelly Guedes de Oliveira Scudine[1], Randriely Merscher Sobreira de Lima[2], Barbara Barth[3], Danusa Mar Arcego[3], Irina Pokhvisneva[3], Zihan Wang[3], Sachin Patel[3], Ciaran Ford[4], Robert Levitan[5], Laurette Dube[3], James Kennedy[5], Roberto Britto Sassi[6], Geoffrey Hal
[1]State University of Campinas, [2]Universidade Federal do Rio Grande do Sul, [3]McGill University, [4]Singapore Institute for Clinical Sciences, National University of Singapore, [5]University of Toronto, [6]McMaster University, [7]Federal University of São Paulo (UN

Obesity and related metabolic disorders such as type 2 diabetes (T2D) arise to a significant extent from the cumulative effect of complex gene-by-environment (G*E) interactions that cumulate over a person's life, impacting both decision making processes and behavior as well as body weight and health outcomes. Yet the multiscale mechanisms at play are little known, especially how gene and brain interact with cumulative exposure to mal/adaptive environmental conditions on behaviors that contribute for the development of metabolic disorders later in life. Since genes do not act in isolation but in concert with other genes in molecular pathways we aimed at developing and applying a novel bioinformatics approach to tracing multiscale dopaminergic networks involved in the response to mal/adaptive environment on food choices and body weight outcomes in early childhood. We created a co-expression-based polygenic risk score (ePRS) that reflects variations in the function of the DRD4 gene network in striatum (Striatum-ePRS-DRD4), and explored its effects on food choices, body weight and brain neuroanatomy in response to variations on environmental stress. We also described whether the GXE interaction pattern reflects the classic and pathology-oriented diathesis-stress model of person-X-environment interaction, or the evolutionary-inspired differential-susceptibility model in 129 Canadian children (62 boys; Caucasian). The analysis

revealed that components from the Striatum-DRD4 gene network associated with response to monoamines and epigenetic processes, including histone H3-K27 methylation specifically in distinct portions of parietal, frontal and temporal cortex are different between children from high versus low adversity. Turning to food choice and BMI, we observed significant interactions between the Striatum-ePRS-DRD4 score and environmental adversity on food intake and BMI, in which more adversity is linked to a general pattern of decreased intake and lower BMI in children with a low striatal DRD4 gene network expression, and increased intake/higher BMI in those with a high striatal DRD4 gene network expression. To identify whether these interactions were consistent with the differential susceptibility model, we used the approach of Roisman and associates, and for several interactions there was evidence for differential susceptibility. The striatal DRD4 gene co-expression network acts as a moderator of early life stress' influence on dietary intake in children. Understanding these complex gene by environment relationships can be informative for the development of strategies to prevent obesity risk and metabolic disease.

N-17: Characterizing habit learning in the human brain using fMRI - a replication study

Rani Gera[1], Elizabeth Tricomi[2], John O?Doherty[3], Tom Schonberg[1]
[1]Tel Aviv University, [2]Rutgers University, [3]California Institute of Technology

The dynamics between goal-directed and habitual action control play a key role in determining behavior, both adaptive and maladaptive. It had been suggested that reinforced actions are at first under goal-directed action control but following multiple repetitions can become habitual, as characterized by automaticity and insensitivity to changes in outcome value. While major progress has been made in understanding the neural substrates of habits in animals, we still lack knowledge on how habits are formed and maintained in the human brain. To date, only one study, yet to be replicated, has demonstrated the behavioral transition from goal-directed to habitual action control in humans and implicated the posterior putamen in the process (Tricomi et al., 2009). Aiming to close this gap and comprehensively characterize the neurobehavioral mechanisms underlying habits at both the individual and group levels, we are in the process of replicating the original Tricomi et al study with a target of a registered n=122 participants. Two groups undergo either short or extensive training on instrumental actions after which behavior is tested for sensitivity to reinforcer devaluation. At this point, prior to COVID-19, we scanned n=17 in the extensive training group and n=10 in the short training group. So far we find no indication of greater habit expression (as indicated by sensitivity to outcome devaluation) following extensive training. Nevertheless, some participants exhibit typical goal-directed behavior while others exhibit typical habitual behavior. Performing similar fMRI analyses to the original study in the 3 day group, we do not find a significant effect in the putamen at this juncture. Nevertheless, it is worth mentioning that the results of several key pre-registered analyses with uncorrected maps implicate the putamen in habitual responding and the anterior caudate in goal-directed behavior. If these preliminary results are confirmed

following the completion of data collection, they would be aligned with the animal literature and support a dual-system account of action control with corresponding neural analogues that are selectively dominant in each type of action control. When the full pre-registered sample of N=61 in each of the groups is complete, we will conduct further planned analyses including studying structural plasticity associated with habit formation (based on diffusion weighted imaging data), and determining whether habit formation can be predicted from changes in functional connectivity measured with resting-state fMRI.

N-18: Replay basis of model-based planning and cognitive map maintenance in humans

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

Model-based decision making, which depends on knowledge of the structure of the environment, enables agents to plan and flexibly adapt to changes in the world. This kind of inference can be implemented via replay of trajectories within cognitive maps, such as those representing paths in a maze. Compressed replay of experience has been shown to be critical for memory encoding and decision making in rodents, but in humans, the role of replay is relatively unknown. In our experiment, participants engaged in a model-based reward learning task (adapted from Doll et al. 2015) which contained two separate and randomly alternating abstract environments. Using magnetoencephalography (MEG) and multivariate analysis techniques, we found evidence for compressed neural replay of paths through the environments. During planning, we found that replay strength reflected the need for model-based inference and positively correlated with two variables related to the potential for higher reward: expected option values and per-trial reward stakes. After feedback, replay of current environment paths was stronger after low reward feedback, potentially reflecting policy change. Further, replay strength of the alternative environment was stronger when current world feedback was less informative. It also scaled with the time since last experiencing the alternate environment, potentially supported memory maintenance. Our results provide novel insight into how rapid replay can support model-based decision making, supporting a view of planning as inference, and demonstrate a dynamic prioritization of inference versus cognitive map maintenance in the human brain.

N-19: All you can choose: The influence of choice exclusivity on the process and experience of decision-making

Xiamin Leng[1], Romy Froemer[1], Amitai Shenhav[1]
[1]Brown University

Objective:
 Choosing a single dish from a menu is typically harder and more anxiety-provoking than choosing from the same options in a buffet. In the lab, we typically study the

former, menu-like choices, where choosing one option precludes choosing any of the others. Popular models of choice likewise typically assume some form of competition between options to account for choice exclusivity. Studying choices that lack (or relax) this exclusivity property could provide novel insights into the range of choices we engage with in our daily lives. Here, we developed a novel task that compares exclusive (menu-like) choices, as typically investigated, to non-exclusive (buffet-like) choices, in which participants have the opportunity to choose more options following their initial choice, and used this task to explore the mechanism of choice exclusivity.
 Method:
 Participants (N=17) viewed sets of four consumer products, and were asked to select their favorite item. They were then allowed to either choose any additional items they liked (non-exclusive choice trials) or not (exclusive choice trials). Choice exclusivity was cued at the start of each trial. At the end of the session, participants rated how conflicted they felt during each of the 160 choices.
 Results:
 Relative to exclusive (menu-like) choices, non-exclusive (buffet-like) choices were faster and led participants to experience less conflict. Remarkably, this increase in response speed was not associated with a significant decrease in choice accuracy - participants were similarly likely to select the most valuable item as their first choice in both conditions. Exclusive and non-exclusive choices were also similarly influenced by the relative difference between chosen and unchosen values. In both conditions, participants were also faster to choose as the overall value of the set increased. However, overall value exerted a stronger influence on response speed for non-exclusive relative to exclusive choices, consistent with a less competitive and more independent (race-like) evidence accumulation process. Unlike exclusive choices, non-exclusive choices also generated less conflict the more valuable the options were.
 Conclusion:
 Our study validates a novel paradigm for studying a greater variety of real-world forms of choice, and our results provide novel insights into the impact of choice exclusivity on the dynamics and subjective experiences associated with decision making. In so doing, we lay the groundwork for new approaches to teasing apart mechanisms that make our choices better from those that make our choices (maybe unnecessarily) hard.

N-20: Attentional dynamics in multi-attribute preferential choices

Xiaoozhi Yang[1], Ian Krajbich[1]
[1]the Ohio State University

When making decisions, how people allocate their attention influences their choices. One empirical finding is that people are more likely to choose the option that they have looked at more, as described by the attentional drift-diffusion model (aDDM; Krajbich et al., 2010). However, options often have multiple attributes. Attention is also thought to govern the relative weighting of those attributes (Roe et al., 2001; Amasino et al. 2019). However, little is known about how these two distinct features of the choice process interact; we still lack a model (and tests of that model) that incorporates both option- and attribute-wise attention. Here, we propose a multi-attribute attentional drift-diffusion model (maaDDM) to account for

attentional discounting of both options and attributes. The model assumes that unattended options and attributes are both discounted. We further hypothesized that the unattended attribute of the unattended item should receive a double discount. To compare variants of the model, we used five eye-tracking datasets (two-alternative, two-attribute preferential tasks) from different choice domains. In particular, two of these datasets comprised choices between 50-50 gambles (Smith & Krajbich, 2018), including a food gamble task (N=44) and a money gamble task (N=36). Other datasets involved decisions between options with meaningfully different attributes, including 1) a mini dictator game with two conditions (Smith & Krajbich, 2018; N = 22), 2) a choice task where subjects made hypothetical choices between clothing items paired with brand labels (Philiastides & Ratcliff, 2013; N=28), 3) a choice task where subjects made choices between snack foods presented with two different nutrition labels (Rramani et al., 2020; N=50). The models were fit to choice and RT distributions using RSTAN combined with RWiener. We found stable option-level and attribute-level attentional discount factors across datasets, though non-fixated options were consistently discounted more than non-fixated attributes. In line with our main hypothesis, subjects generally showed a double discount for the non-fixated attribute of the non-fixated option (best model in 6/7 datasets). We found a bi-directional relationship between gaze and attribute weights, with more important attributes receiving more gaze, but also more weight on attributes when they received more gaze. This indicates that when attributes are meaningfully different, additional attribute weights are needed to fully account for behavior. In summary, our work uncovers an intricate interplay between attribute weights, gaze processes, and preferential choice.

N-21: Attentional Sampling Strategies in Multi-Attribute Decision Making

Jacob Elsey[1], Ernst Niebur[1], Veit Stuphorn[1]
[1]Johns Hopkins University

Many everyday decisions require us to consider multiple options, each comprised of a number of shared or different attributes. In such situations, optimal decisions require the simultaneous integration and comparison of a large amount of information. This complexity is thought to tax limited cognitive resources, requiring the use of attention to select a limited subset of information, on which the decision is then based. To better understand the role of attention in decision making, subjects participated in multi-attribute decision making tasks in which we systematically varied the complexity of choices by increasing the number of options, or attributes, or both. Our task design uses eye position as an overt measure of attention while subjects choose between gambles that systematically cover the decision space. Each attribute's magnitude is masked by its corresponding color cue, and only revealed when fixated. Subjects freely inspect the attributes with no time constraint, before indicating their choice with a keypress. The resulting eye movements provide temporal and spatial information about subjects' focus of attention during the attribute sampling process. We find that human participants mainly switch between two strategies for sampling decision-related information.

When only few options are available, participants most frequently use an exhaustive strategy to sample all attributes of each option independently, before examining the next one (2-Option, 2-Attribute: 68% of subjects, 2-Option, 4-Attribute: 66%). This is an efficient strategy because all information is examined. However, when more options are available, participants most commonly employ a filtering strategy. They first compare their most important attribute across all options to identify the one or two best ones, and then sample all attributes of the remaining options (4-Option, 2-Attribute: 75% of subjects, 4-Option, 4-Attribute: 68%). This strategy reduces the amount of information that needs to be processed but can lead to suboptimal choices. Importantly, which strategy is used depends on the number of options, but not the number of attributes. Thus, increasingly complex decision problems evoke strategic shifts in the direction of attention during the search for the preferred option. Because our new task design provides detailed information over the values and temporal order of sampling the attributes, it retains the ability to control task requirements typical for laboratory designs, while still capturing important aspects of real world decision making in complex environments. Thus, this task allows us to explore the neuronal mechanisms underlying the different strategies.

N-23: Efficient coding of numbers explains decision bias and noise

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

Multi-attribute decisions require the encoding of several pieces of information, followed by their aggregation. A simple example is the computation of an average (such as the average return on a risky investment). When asked to compare averages of numbers, human subjects have been reported to differentially weight different numbers, even though these should be equally relevant to a correct decision. This selective weighting has been interpreted as resulting from a biased encoding of the numbers, that optimally compensates for the presence of internal noise arising later in the decision process (Spitzer, Waschke, and Summerfield, Nature Human Behavior, 2017). A natural alternative assumption is that numbers are encoded with noise, then optimally decoded from the noisy representations. Under the hypothesis of efficient coding, the degree of encoding noise should vary in an optimal way across the stimulus space, and should depend on the statistics of the numbers. We investigate these predictions through a task in which subjects are asked to compare the averages of two series of numbers, each sampled from the same prior distribution. We manipulate the shape of this prior distribution across blocks of trials (it can be either uniform or skewed towards smaller or larger numbers). We find that in addition to encoding each number with a bias, our subjects seem to encode the numbers with a degree of noise that varies with the size of the number. In particular, numbers that are unlikely under the prior are encoded with greater noise. We further show that a maximum-likelihood decoding model captures subjects' behavior, and indicates that the encoding is efficiently adapted to the prior, resulting in higher expected rewards in the task. Furthermore, our model predicts a relation between the bias and the variability of estimates,

thus providing a statistically-founded, parsimonious derivation of Wei and Stocker's "law of human perception" (Wei and Stocker, PNAS, 2017). Our results both shed new experimental light on the mechanisms by which human subjects process numerical information, and increase our theoretical understanding of an encoding-decoding paradigm in which efficient coding is combined with inference from noisy internal representations, resulting in biases and variability at the behavioral level.

N-24: Exploring molecular contributions to disease-relevant alterations in reward processing: Disruption of Nrxn1a within excitatory forebrain circuits drives value-based dysfunction.

Marc Fuccillo[1], Joe Kable[1], Opeyemi Alabi[1]
[1]University of Pennsylvania

Background: Reward learning is essential for daily function and impaired across neuropsychiatric disorders. While neural circuit mechanisms mediating these behaviors are receiving increased attention, how neuropsychiatric disease-associated mutations disrupt these systems remains largely unexplored. Neurexin1α (Nrxn1α) is a presynaptically-localized molecule with major, reproduceable associations to autism and schizophrenia, disorders exhibiting significant reward-learning abnormalities. Nrxn1α is expressed throughout the mouse brain but enriched in key cortico-striato-thalamic nodes. Here we examined mice with mutations in Neurexin1α in value-based choice paradigms with 3 goals - (1) quantify choice deficits in Nrxn1α genetic models; (2) localize these behavioral deficits to Nrxn1α disruption in specific brain circuits; (3) describe changes in value-based neural signals following circuit-specific Nrxn1α perturbation. Methods: Exp.1. Male Nrxn1α-knockouts (KOs, missing brain-wide Nrxn1α) (n=10) and wildtype (WT) littermates (n=11) were run in feedback-based operant paradigms where mice had to choose between two alternatives varying in outcome benefit and required operant effort, while dynamically tracking changes in the cost-benefit contingencies across time. Exp.2. Deletion of Nrxn1α was restricted to developing cortical excitatory neurons (Nrxn1α-cortKOs, n=13) or thalamic excitatory neurons (Nrxn1α-thalKOs, n=10), and these mice were compared to their respective littermate controls (n=11/8, respectively) in value-based operant tasks. Exp.3. Direct pathway striatal spiny projection neurons (dSPNs) were imaged via fiber photometry of genetically-encoded calcium indicators in Nrxn1α-cortKOs (n=7) and controls (n=6) during operant behavior. Results: Nrxn1α-KOs exhibit blunted choice towards outcomes associated with greater benefits ($p<0.05$) or fewer costs ($p<0.05$). Reinforcement learning models suggested this was driven by deficits in the updating and representation of choice values (both $p<0.05$). Neurexin1α disruption within cortical excitatory projection neurons ($p<0.005$), but not thalamic populations (ns), recapitulated most aspects of the whole-brain KO phenotype. In vivo striatal dSPN population calcium activity demonstrated that cortical-selective Neurexin1α KO disrupts reward-associated neural signals within striatum ($p<0.05$), a key site of feedback-based learning. Conclusions: By relating deficits in value-based decision-making to region-specific Nrxn1α disruption and changes in reward-associated neural

activity, we highlight potential neural substrates for pathophysiology of neuropsychiatric disease-associated cognitive dysfunction.

N-25: Neuroscience, psychiatry & big data: Using neuroeconomic risk/reward metrics to predict destructive behavior

Leandros Stefanopoulos[1], Byoungwoo Kim[2], Khrystyna Stetsiv[2], Sumra Bari[2], Nicole Vike[2], Nicos Maglaveras[1], Hans Breiter[2], Aggelos Katsaggelos[2]
[1]Aristotle University of Thessaloniki, [2]Northwestern University

Introduction: Recent increases in computing power have allowed acquisition of large datasets on consumption, and other human behaviors related to mental health (MH) and destructive behavior (DB) (e.g., suicidality and antisocial action). These big data approaches offer the opportunity to transform the way we research neuroeconomics to quantify human preference behavior and make predictions about health and disease. To date, approaches using text mining, self-reports, and portable sensors have had limited capacity to predict DB/MH. Here, we investigated whether a relative preference theory (RPT) framework characterizing approach/avoidance behavior can be used to predict DB and/or MH at the level of hundreds of people. Methods: Subjects (N=501) rated pictures from the International Affective Picture System (IAPS) from multiple categories (e.g., nature scenes, aggressive animals, men/women in bathing suits, etc.) on a scale from -3 to 3. They also rated 21 recent and current categories of DB/MH of a 7-point Likert scale. Measures of the standard deviation (σ) and Shannon entropy (H) of ratings within each picture category were plotted these metrics against the mean rating (K) for each category. Fifteen features from these RPT graphs (e.g., loss aversion, risk aversion) were used in multivariate logistic regression to predict each of the 21 DB/MH and quantify mean absolute errors (MAE). Results: We report four principal findings: (1) When H was plotted against K for each picture category, a value function emerged that closely resembles the value function established for prospect theory by Kahneman & Tversky. (2) When plotting rating σ against K, we observed a saturation function in which σ rises to a peak value for intermediate values of K before declining back to smaller values as K increases further. These saturation dynamics paralleled the decision utility from portfolio theory as elaborated by Markowitz. (3) RPT graph features predicted DB for "being disruptive", "desire to start fires", "breaking rules" and "wanting to hurt yourself" with MAE of 0.2 - 0.47. (4) RPT graph features predicted MH for "intrusive thoughts", "paranoia", "sadness" and "intense fear" with MAE of 0.5 - 1.0. Conclusions: (1) RPT extends to human approach/avoidance behavior measured through ratings, validating previous work with operant keypress tasks. (2) DB prediction, including suicidality, from RPT features was very strong, and stronger than prediction of MH. Together these findings argue that neuroeconomic features combined with multivariate classification have utility for predicting DB, and suggest one route how neuroeconomics can impact AI and psychiatry.

N-26: Investigating individual differences in latent structure learning in a changing environment

Evan Dastin-van Rijn[1], Linda Yu[1], Alexandre Filipowicz[2], Zhiyi Chen[3], Tingyong Feng[3], Matthew Nassar[1]
[1]Brown University, [2]Universtiy of Pennsylvania, [3]Southwest University

Objective: Whether faced with the volatility of the stock market or the availability of goods on grocery store shelves during the pandemic, it is important to learn adaptively in changing environments. Critical to successful adaptive learning is the ability to identify the relevant causal aspects of the environment (known as latent structure) in order to figure out which changes are enduring and which are momentary. For example, is that dip in the stock market indicative of a true crash, or just the latest jitters from the news cycle? Is the shortage of toilet paper indicative of a breakdown in the supply chain, or is the local store just temporarily in between restocking? Individuals exhibit great variability in how much they change behavior in response to deviations from expectations, and in how well they recognize latent structure in tasks (Nassar & Trioani, 2020; Collins & Koechlin, 2012; Vaidya et al., 2020). Which characteristics underlie these individual differences are not well known, but identifying them would be beneficial for policy, business, and health interventions. The objective of this study is to characterize the neural and behavioral architecture of individual differences in latent structure and adaptive learning in a large sample of participants. Methods: We analyze behavioral and resting state fMRI data in 294 individuals who performed a predictive learning task. In the task, participants have to infer the locations of two hidden helicopters, based on the gold bags that are dropped in Gaussian distributions around each (the helicopters trade off on dropping bags). Participant behavior is fit with a weighted linear combination of a standard Rescorla-Wagner (RW) and a Bayesian ideal observer (IO) model in order to capture the degree to which they utilize latent learning. We also correlate the relative weight of latent learning from this hybrid model with individual differences in resting state functional connectivity. Results: For the majority of the participants, the hybrid model was found to fit behavior better than either a RW or IO model alone. A median split on the weight of combination showed that participants fit by a more IO-like model had updates that were significantly more influenced by the true position of the helicopter as the task progressed (2 sample t-test, $p<10e-6$). Ongoing analyses seek to identify subnetworks that are strengthened in individuals whose behavior is most influenced by latent states. Conclusion: Participants fall along a continuum in their usage of RW vs. latent learning strategies, and their relative weights of latent learning describe their success on learning the task structure over time.

N-27: Adaptive attentional shifts amplify the framing effect under time constraints

Ian Roberts[1], Yi Yang Teoh[1], Cendri Hutcherson[1]
[1]University of Toronto

Dual-system models of choice assert that rapid, automatic processes (i.e., System 1) facilitate quick action, but sometimes produce seemingly irrational choice biases that must be corrected by slower, more deliberative processing (i.e., System 2). In these models, time pressure can amplify choice biases because it disrupts System 2 while leaving System 1 processing intact. For example, recent work suggests that time pressure amplifies framing effects because it prevents the activation of slower but more rational preferences, allowing quick, framing-sensitive processes to drive decisions. Here, we develop an alternative account of these effects, situated within a model of rational attention. Specifically, we assume that in many decision problems, information is gathered serially, and that attention is thus deployed in a strategic manner to maximize value. This model predicts that time constraints might alter information search, leading to changes in choice biases. To investigate this attentional mechanism, we conducted a preregistered replication of a recent set of studies where time pressure amplified the framing effect (Guo et al., 2017) while recording participants' eye movements (N=40). As predicted, time constraints induced strong attentional shifts that accounted for previously-observed behavioral effects. We find that attention under time constraints contributed to the increased framing effect via the combined influence of two mechanisms: (1) increased influence of peripheral cues predictive of reward on first fixation and (2) increased probability of choosing after only fixating one option. Building on these findings, we propose a novel computational model where value-guided information sampling via attention accounts for the impact of time pressure on the framing effect. A rational attention model thus provides novel insights into when and why time pressure amplifies choice biases while replicating behaviors previously attributed to dual-systems. It raises important questions about the extent to which systematic choice biases such as the framing effect may be understood as emergent properties of an interaction between goal-directed decision strategies and attentional processing constraints.

N-28: Neuropharmacology of approach and avoidance: Disentangling value and salience processing

Geraldine Gvozdanovic[1], Ana Cubillo[1], Robert Schreiber[1], Lydia Hellrung[1], Thorsten Kahnt, Boris Quednow[1], Philippe Tobler[1]
[1]University of Zurich

Exerting adequate approach and avoidance behavior is essential for survival. Value and salience are two fundamental concepts that influence approach and avoidance behavior but have often been confounded in previous research on value-based decision making. While previous research associated both value and salience with dopaminergic functioning, recent studies propose that norepinephrine contributes to processing motivational salience. Here, we investigated two types of salience, elemental and global salience, which differ when outcomes

comprise both gains and losses. Global salience corresponds to absolute expected value, whereas elemental salience sums up the absolute values of each element. We administered preferential norepinephrine (Reboxetine) and dopamine reuptake inhibitors (Methylphenidate) before participants (n = 118) performed a value-based decision task. During a learning phase, participants first associated geometric figures and colors with different monetary outcomes (+/- 1 CHF, +/-5 CHF and +/- 10 CHF). In the subsequent decision phase, participants either accepted or rejected combinations of the learned colors and figures. Response times were measured for each trial, individual averages regressed against expected value, elemental salience and global salience and regression weights compared against zero (one-sample t-tests) or between decision types (paired t-tests). Overall, our results indicate that expected value ($t_{117} = -2.462$, $p < 0.01$), elemental salience ($t_{117} = -7.667$, $p < 0.001$), and global salience ($t_{117} = -6.078$, $p < 0.001$) each accelerate response times during decision making. Elemental salience had a stronger accelerating effect on reaction times than global salience ($t_{117} = -1.810$, $p < 0.05$). With regard to pharmacology, we predict norepinephrine reuptake blockade to increase primarily the accelerating effect of salience on reaction times. Conversely, we predict preferential dopamine reuptake blockade to primarily increase the accelerating effect of expected value on response times. The findings will have implications for altered approach and avoidance behavior in neuropsychiatric disorders.

N-29: Dynamic changes in feedback valuation explain learning in a probabilistic two-armed bandit task

Shabnam Hakimi[1], Benjamin Muzekari[1], Kelly Eom[1], Sonakchhi Shrestha[1], Candice Wang[1], Stephanie Ng[1], Alannah Rivera-Cancel[1], John Thorp[1], Savannah Erwin[1], Rachael Wright[1], Nancy Zucker[1], Alison Adcock[1]
[1]Duke University

Realizing the outcomes of our actions and using this feedback to optimize future behavior is central to learning. Whether individual learners actually value this feedback is less clear. We developed a novel probabilistic two-armed bandit task where individuals could elect to work to reveal obscured bandit choice feedback, thus eliciting a dynamic valuation signal for feedback in learning. Here, we demonstrate the efficacy of this approach for preference elicitation and investigate the relationship between feedback valuation and task performance. Over 160 trials, 101 Duke undergraduates chose between two probabilistic bandits (pwin=0.8 or 0.6) and received either veridical feedback (0 or 1) or non-informative, neutral feedback (centered on 60% or 40% of trials, respectively). In half of the non-informative feedback trials (20% overall), participants could choose to work for feedback (yes/no). Selecting 'yes' unlocked a 2.5 s-period where each spacebar press was linked to an incremental increase in the transparency of a shape occluding the feedback cue; participants were instructed to respond as quickly as possible to reveal feedback. For each participant, we computed moving averages (window size=16) over bandit choice, choice to work for feedback, and the local availability of the work-

for-feedback choice (ratio of work-for-feedback to veridical feedback trials). We also computed dOptimal, the relative preference for the better bandit over time. We then fit two nonlinear mixed effects models to the data to determine whether the choice to work for feedback and its relative salience predicted learning to choose the better bandit. M1 tested the relationship between dOptimal and feedback choice over time ($p=0.12$). M2 also accounted for the salience of the feedback choice and better explained the data (AIC & BIC: M2 < M1). Learning over time was dynamically modulated by the choice to work for feedback ($F_{1,12525}=3.717$, $p=0.053$), although this relationship was statistically significant only when accounting for local changes in work-for-feedback salience (interaction term $F_{1,12525}=11.81$, $p<0.001$). Work-for-feedback salience also predicted learning over time ($F_{1,12525}=13.29$, $p<0.001$), reflecting the impact of environmental statistics on learning. Using this novel approach, we demonstrate that individuals are willing to work to reduce uncertainty about feedback during learning and that individual preference for feedback is related to task performance. We also show that feedback valuation changes dynamically with environmental statistics and that these shifts can significantly impact the trajectory of learning.

N-30: Neurogenetic Evidence for Sex Differences in Neural Pathways Underlying Food Cue Responses in Adolescents: Prefrontal DRD4 Gene Expression Effects

Afroditi Papantoni[1], Andre Portella[2], Liuyi Chen[3], Richard Lee[3], Robert Levitan[4], Patricia Silveira[2], Laurette Dube[2], Susan Carnell[3]
[1]University of North Carolina at Chapel Hill, [2]McGill University, [3]Johns Hopkins University School of Medicine, [4]University of Toronto

Obesity has been associated with increased impulsive decision-making and attention bias in response to food cues in the environment, and there is evidence that the mesocorticolimbic dopamine signaling pathway plays a key role in these executive functions. Previous studies have linked hypofunctional polymorphisms of the dopamine receptor 4 gene (DRD4) to reduced inhibitory feedback and increased sensitivity to rewarding cues. Adolescence is a critical period for both obesity and brain development, with striking sex differences. We examined how functional brain activation in response to food and non-food cues is moderated by the prefrontal predicted DRD4 gene expression and sex. 76 adolescents (37 boys and 39 girls, age 14-18y, mean BMIz 0.54 SD 1.23) underwent an fMRI to assess BOLD activation in response to pictures of high energy-density (high ED) foods, low energy-density (low ED) foods and non-food objects, in a fed and fasted condition. Whole brain group-level analyses, using Bayesian inference, were applied to contrast images generated from first-level analyses of brain response activation to high ED vs. low ED food cues and food vs. non-food cues (effect size $d>0.2$, $\log BF > 3$). Genotyping data obtained from saliva DNA was used to estimate individual levels of predicted prefrontal DRD4 gene expression using PrediXcan. Adolescents were divided into two groups (mean split) exhibiting relatively low vs. high DRD4 expression. Our results showed a distinct pattern of activation for boys and girls. In the fed condition, girls with lower

expression of DRD4 in the prefrontal cortex had lesser BOLD activation in the orbitofrontal cortex, premotor cortex and lingual gyrus for low ED foods (vs. non-foods), while boys with lower expression levels had greater BOLD activation in the inferior parietal lobule for both high ED and low ED foods (vs. non-foods). In the fasted condition, girls with lower (vs. higher) expression of DRD4 in the prefrontal cortex had greater BOLD activation in the putamen and inferior frontal gyrus for high ED (vs. non-foods), while boys with lower expression had greater BOLD activation in the inferior frontal gyrus and lingual gyrus for high ED foods (vs. non-foods) and greater BOLD activation in the dorsolateral prefrontal cortex and cerebellum for low ED (vs. non-foods). Our findings suggest that prefrontal DRD4 gene expression may influence responses to food cues via distinct neural pathways in adolescent boys and girls. Further research is needed to replicate these effects, explore relationships with eating behavior, and investigate impacts of prefrontal DRD4 expression on connectivity within appetitive circuits.

N-31: Habitual behavior in humans can be experimentally induced by varying the length of training on an instrumental action.

Aniek Fransen[1], Omar Perez[1], John O'Doherty[1]
[1]California Institute of Techology

In rodents, extended training on an instrumental action has been found to render behavior habitual, e.g., it becomes automatic and independent of the subjective utility of the outcome produced by the action. Despite the theoretical importance of habits for understanding human behavior, the literature on the experimental induction of habits in humans remains weak and inconsistent. Here, we report the successful development of a new paradigm that prompts habitual control in humans by means of the manipulation of the training duration. We developed a paradigm in which participants learn to perform two instrumental actions to obtain one of two different token outcomes that have equivalent monetary value. After the training phase, one of the tokens is devalued. In a final choice test between the two previously-trained actions, participants are asked to perform either of the actions in pseudo-extinction (i.e., outcomes are not presented). We varied the length of instrumental training the participants were exposed to and tested for a difference in the number of times the now devalued action is selected during the choice test between the groups. Across two experiments, acquired in-lab (N = 58) and through Amazon's Mechanical Turk (N = 412), the group exposed to more extensive training became significantly less sensitive to devaluation, suggesting habitual behavior. In the larger MTurk sample we also found that the expression of habitual behavior after extended training is most pronounced at the beginning of the pseudo-extinction test phase; a sign that our extensively trained participants relied more on habitual control initially but reverted back to goal-directed control over the course of the test phase. By contrast, the behavior of participants who only received a limited amount of training exhibited goal-directed control throughout the pseudo-extinction phase. We show in two separate experiments that extended training on an instrumental action can render human responding habitual. These

results provide the clearest evidence to date that humans do indeed become increasingly habitual the more an instrumental action is performed and repeatedly reinforced. Humans, just like animals, are not exclusively goal-directed, but instead develop habitual behavior as a function of experience. The development of a reliable means to induce habits in humans is a crucial prerequisite toward the delineation of the neural mechanisms of habitual control in the human brain, as well as for facilitating a more meaningful characterization of how those mechanisms can become altered in disease.

N-32: The relationship between positive emotions, coping behaviors and foraging decisions during the COVID-19 pandemic

Vera Ludwig[1], Vani Gupta[1], Arjun Ramakrishnan, Michael Platt[1]
[1]University of Pennsylvania

The COVID-19 pandemic has taken severe tolls on mental health. While its effects on negative mental states, e.g. anxiety, have been widely documented, less is known about the role of positive emotions during the pandemic, and about how a presumed lack of positivity affects decision-making. Specifically, a lack of positive emotions during the pandemic may be linked to reduced exploratory behavior or suboptimal choices. We investigated the relation between positive emotions, coping behaviors and explore-vs-exploit decisions during the pandemic using a virtual foraging task. The task involved collecting berries at different patches. Berry harvest rates per patch declined over time, and participants could decide to switch to a new patch at any time. Moving to a new patch took time, creating opportunity costs. Previously, we showed that anxiety is linked with suboptimal choices on this task. However, the relation of the task with positive emotions has not yet been studied. We hypothesized that certain coping behaviors during the pandemic (e.g., engaging in fun activities) predict positive emotions and that such positivity systematically affects the tendency to explore or exploit, above and beyond the effects of negative emotions (state anxiety). We obtained valid data from 483 participants in a cross-cultural sample (240 from a university in India, 209 via MTurk in the US, 34 via the Wharton Behavioral Lab, US). Participants completed the PERMA-profiler, including a positivity score, and the Spielberger's State-Trait anxiety inventory. Participants also reported the daily number of hours spent on twenty activities, including chores, socializing, etc, and played two blocks of the foraging game (~10 min). Participants were instructed to pick as many berries as possible and were monetarily rewarded for their performance. Participants reporting to spend more time than others doing fun and pleasurable activities also reported to feel more positive overall. As predicted, foraging was systematically linked with this positivity, as shown by a significant interaction effect between the time spent on a patch and positivity when predicting the likelihood of participants switching patches. For highly positive participants, the time spent on a patch predicted switching less strongly than it did for less positive participants. These results held up after controlling for state anxiety. Our results indicate that it is worthwhile studying the effects of both positive and negative emotions on decision-making. Moreover,

they highlight the roles of various coping behaviors to foster positivity and optimal decision-making during high adversity such as the current pandemic.

N-38: The structure of white-matter connections from the anterior insula to the nucleus accumbens predicts relapse to stimulant drug use

Loreen Tisdall[1], Kelly MacNiven[1], Josiah Leong[2], Brian Knutson[1]
[1]Stanford University, [2]University of Arkansas

Nucleus Accumbens (NAcc) activity can predict economic and addictive choice. Previously, we found that the structural coherence of ascending dopaminergic midbrain projections to the NAcc was related to diagnosis of Stimulant Use Disorder (SUD), but not relapse (MacNiven et al., 2020). Here, we further tested whether the structural coherence of glutamatergic connections that project to and modulate activity in the NAcc might instead predict relapse to stimulant use at 6 months. Using Diffusion-Weighted Imaging, we identified and assessed structural properties of white-matter tracts connecting the Medial PreFrontal Cortex (MPFC), Anterior Insula (AIns), and Amygdala (Amy) to the NAcc. To address methodological challenges for identification and quantification of the tracts (e.g., crossing fibers and inter-subject anatomical variability), we fit a constrained spherical deconvolution model and conducted probabilistic tractography in individuals' native space. We then tested if the tracts' structural properties (i.e., fractional anisotropy (FA) and inverse mean diffusivity (1-MD)) could predict relapse 6 months after treatment release. Probabilistic tractography successfully resolved all of the white-matter tracts in 110 participants, including 40 healthy controls (57% male, 32.7±11.8 years old) and 70 patients with SUD (81% male, 41.4±11.7 years old). Regarding treatment outcome at 6 months, we obtained a comparable sample of individuals who had abstained (n = 30, 80% male, 39.5±12.5 years old) and who had relapsed (n = 28, 86% male, 42.3±10.9 years old); 12 SUD patients (75% male, 43.8±11.9 years old) were lost to follow-up. Controlling for age, sex, and head motion, lower tract coherence (FA, 1-MD) in the right AIns tract predicted relapse (FA z=-2.059, p=0.039; 1-MD z=-2.019, p=0.044). Additionally, lower left Amy 1-MD predicted relapse (z=-2.19, p=0.029). Structural properties of the MPFC-NAcc tract, however, did not predict relapse (although they did decline with age). None of these coherence metrics were associated with SUD diagnosis. Together, our findings suggest that the structural properties of specific white-matter tracts are selectively associated with relapse to stimulant use, and identify novel targets for prediction of relapse. While the study did not incorporate functional activity, the localization of structural projections from the AIns to the NAcc might implicate drug cue-induced negative arousal, rather than value integration commonly associated with MPFC-NAcc projections in relapse. Future research may test these predictions by linking structural tract properties to functional brain activity and risky behavior in the context of relapse.

N-39: Motivated performance while sleep deprived: reduced ACC and insula recruitment and increased effort-discounting.

Stijn Massar[1], Julian Lim[1], Karen Sasmita[2], Michael Chee[1]
[1]National University of Singapore, [2]Cornell University

Sleep deprivation (SD) has a negative impact on the motivation to exert effort. This may contribute to the decline in attentional performance observed under SD. In this study we examined how SD affects motivated performance and effort-based decision making. Particularly, we used functional magnetic resonance imaging (fMRI) to uncover the neural mechanisms underlying the interplay between SD and motivated behavior. Twenty-seven healthy subjects were tested once after a night of sleep in the lab (Rested Wakefulness = RW), and once after a night of total sleep deprivation. Participants performed an effortful attention task with different incentive levels (0, 10, or 50 cents/fast and correct response). Behavioral performance and fMRI data were collected during task performance. Subsequently, participants performed an effort-based choice task, during which they could choose to earn additional rewards for performing the attention task for a longer duration. As expected, attentional performance was worse in the SD session compared to the RW session. In addition, performance improved as a function of incentive level both in the RW and the SD session. This reward-effect was accompanied by increased activation in attention-related brain areas (including dlPFC, ACC/dmPFC, and anterior insula), and increased arousal-related thalamus activation. This reward-modulation was more extensive during RW than SD, particularly in the anterior cingulate cortex (ACC) and anterior insula (aIns; Areas involved in effort regulation). Results from the decision making task showed that participants discounted reward value based on the proposed duration of the attentional task in both the RW and SD conditions. Computational modeling demonstrated that choice behavior was best fit by a sigmoid discounting function. Both model-free and model-based analyses indicated that participants discounted more strongly during the SD session compared to the RW session. In conclusion, our data show that performance decline after SD is reward-dependent and willingness to perform is reduced. Reward modulation of attention-related brain activation is reduced, particularly in areas that are associated with effort regulation (i.e. ACC and aIns). These findings indicate that motivational factors contribute to decline of vigilance following sleep deprivation.

N-40: The neural plasticity of creativity learning: An intervention of mindful learning

Yu-chu Yeh[1], Chih-Yen Chang[1], Chih-Yen Chang[1]
[1]National Chengchi University

Mindful learning is a mindful learning process in which individuals actively and consciously pay attention to the things they are curious about or interested in, and further, try to bring about new meanings or original thinking from these ordinary or special things; it is a way to balance brain networks and optimize the creative process. However, no study has been performed to

examine the neural plasticity with regard to creativity after mindful learning interventions. In this study, we conducted a pretest-posttest design to investigate neurobiological correlation of creativity before and after a mindful learning intervention through a fMRI scanner. Participants were 20 college students who were requested to mindfully take photos using mobile phones and share the imaginative narratives about the photos they took online during the 10-day intervention. The results showed that mindful learning was associated with better creativity scores as well as increased brain activation in bilateral dorsal and ventral anterior cingulate cortex. Dorsal anterior cingulate cortex involves a top-down control mechanism that contributes to the process of evaluating generated creative ideas. Highly activation of this area may be related to inhabiting the stereotypical thinking in order to explore an original idea. In addition, ventral part of anterior cingulate cortex plays an important role on emotional regulation and self-control. The findings of this study suggest the employed mindful learning intervention contributes to facilitating creativity by enhancing the inhibition of non-creative ideas, self control, and emotional regulation.

R-1: A gut feeling: how your brain (and your guts) defines your choices

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

Risk-taking behavior and intertemporal choices are common sources of inconsistences in decision-making. Although both phenomena have high economic and societal impacts, little is known about their underlying neurological processes. This knowledge is essential to understand its mechanisms and correlations. Since this remains an open question in economics and neuroscience, we explore such underlying mechanisms, aiming to clarify if these processes are correlated and how they are processed. Risk-taking behavior was chosen as our topic based on the positive correlation between a rich gut microbiota and risk-taking behavior with rodents (Chumney & Robinson, 2006). Considering the promising results presented by studies replicating such experiments with humans (Dinan et al., 2013), we believe we can prove the influence of the gut-brain axis on risk-taking behavior. Moreover, based on the possible common neural mechanisms underlying risk and time preferences (Bechara & Damasio, 2005), we also investigate the gut-brain influence on intertemporal choices. To explore such neurological processes we manipulate the gut microbiota, employing our most innovative contribution, which is the use of probiotics to investigate how the gut-brain axis is involved in decision-making. Our experiment uses a mixed design with two sessions (28 days gap), and different economic games to elicit risk-taking behavior and intertemporal choice. The preliminary results indicate that risk-taking and intertemporal choices are only correlated in specific scenarios. This correlation was present when participants faced risk-choices with low contrast between the options' expected values, and with intertemporal choices involving longer intervals. The probiotics intervention did not affect participants' stress response or intertemporal choices. However, a decrease in risk-taking behavior was observed after

probiotics administration, which suggests an effect of the microbiota-gut-brain axis on decision-making under risk. Keywords: Gut-Brain, Probiotics, Risk, Intertemporal choice

R-2: How suggestion affects cognitive regulation, and through it hunger perception.

Solène Frileux[1], Solène Frileux[2], Liane Schmidt[1]
[1]Brain and Spine Institute, [2]Pitié Salpêtrière Hospital

Background Research across domains indicates that suggestion about the efficiency of a consumer good or a medical treatment is underpinned by activation of frontal cortex regions associated to cognitive regulation1,2. Though, the recruitment of mechanisms of cognitive regulation such as the propensity to weight in the attributes of a treatment context, while evaluating a symptom under the suggestion of clinical benefit, has never been shown. Methods Here we used hunger evaluations in 126 healthy participants as a model, in order to validate a paradigm for a future clinical study. We hypothesized that (1) suggestions about appetite should influence hunger ratings after fasting, and (2) this effect should be mediated by the participant's regulatory success. Regulatory success was defined by the difference in accepting to eat [unhealthy, tasty] versus [healthy, untasty] food during a validated food choice paradigm from decision-neuroscience4. We used mediation analyses to test our hypotheses. Results We found that appetite suggestion significantly influenced hunger ratings (total path c, r=0.75, p=0.01). Participants under the diminished-appetite suggestion displayed more regulatory success than participants under the enhanced-appetite suggestion (path a, r=-0.11, p=0.03), and participants with more regulatory success reported lesser hunger irrespective of suggestion (path b, r=-1.69, p=0.001). Importantly, the product of the two univariate effects, a and b, significantly mediated the direct effect of appetite suggestion on hunger ratings (path a*b, r=0.18, p=0.02). Conclusion Taken together, our results showed that suggestion about appetite generated cognitive regulation processes that lead participants to perceive less hunger. These results provide insights into how suggestions get translated into an effect on behaviour, opening the window toward a better understanding of the psychological processes that contribute to placebo effects. 1. Benedetti, F. et al. Neurobiological mechanisms of the placebo effect. J. Neurosci. Off. J. Soc. Neurosci. 25, 10390-10402 (2005). 2. Wager, T. D. & Atlas, L. Y. The neuroscience of placebo effects: connecting context, learning and health. Nat. Rev. Neurosci. 16, 403-418 (2015). 3. Hutcherson, C. A., Plassmann, H., Gross, J. J. & Rangel, A. Cognitive regulation during decision making shifts behavioral control between ventromedial and dorsolateral prefrontal value systems. J. Neurosci. Off. J. Soc. Neurosci. 32, 13543-13554 (2012). 4. Harris, A., Hare, T. & Rangel, A. Temporally dissociable mechanisms of self-control: early attentional filtering versus late valulue modulation. J. Neurosci. Off. 33, 18917-18931 (2013)

R-3: Cortical oscillatory changes during an effortful Go/NoGo task in relation to expected monetary gains and losses.

Adam Byrne[1], Danielle Hewitt[2], Jessica Henderson[2], Alice Newton-Fenner[2], Hannah Roberts[2], John Tyson-Carr[2], Nicholas Fallon[2], Timo Giesbrecht[3], Andrej Stancak[2]
[1]MrUniversity of Liverpool, [2]University of Liverpool, [3]Unilever

Aims. Losses have been found to have greater subjective value than gains of equal nominal value in economic decision-making tasks, but do not cause a relative improvement in performance. We explored how the subjective value of losses and gains affected performance and oscillatory changes over sensorimotor electrodes in the alpha and beta bands during a cued Go/NoGo task. Methods. Participants completed the Go/NoGo task while expecting a reward (+10p), avoiding a loss (-10p), or receiving no incentive (0p) if they responded faster than their median RT to a Go cue, or successfully withheld their response to a NoGo cue. Pre-movement alpha- and beta-band EEG power was analysed using the event-related desynchronization (ERD) method, and the subjective value of effort was analysed using a cognitive effort discounting task (COGED). Results. RTs were faster and NoGo error rates were higher when participants expected a monetary gain. Bilateral sensorimotor alpha and beta ERD was stronger when participants were offered an incentive compared to no incentive, and frontal alpha ERD was stronger when participants expected a monetary gain compared to both a loss and no incentive. The COGED task revealed that participants were more willing to expend a greater effort when they were motivated with losses compared to gains, however the value of effort was not associated with ERD patterns. Conclusions. Results show that, while losses are more motivating than gains, positive motivation favours both performance and anticipatory attention during an effortful task which required both motor activation and inhibition.

R-4: Attention and Decision Making in Complex Risky Choice

Sudeep Bhatia[1], Lisheng He[1]
[1]University of Pennsylvania

Prominent cognitive and neural theories of choice, such as the attentional drift diffusion model (ADDM), assume that preferences accumulate in additive increments, guided by attention to the attributes of choice options. Yet in order to make reasonable decisions, the attributes of options often need to be aggregated in a multiplicative or interactive manner. In risky choice, for example, subjective expected utility theories (SEUT) propose that payoffs are multiplied against probabilities. This allows decision makers to successfully evaluate multi-branch gambles and avoid violations of simple dominance. It is unclear how such a mechanism could be implemented within additive models such as ADDM. This is a type of binding problem, one that is closely related to the problem of combing the features of objects in perceptual tasks. We present a solution to this problem in the domain of risky choice. Specifically, we propose that decision makers do accumulate preferences in additive increments guided by attention.

However, attention displays nuanced interactive dynamics, with attention to high probability values increasing subsequent attention to payoff values within the same branch of the gamble. We show that these dynamics generate accumulation rates that are multiplicative functions of gamble probabilities and payoffs, mimicking the desirable normative and descriptive properties of SEUT. We formalize this mechanism with a Markov process for attention, combined with an ADDM accumulator for aggregating sampled attributes. We test the model using Bayesian model fitting on the individual-level in four existing and two new eye-tracking and mouse-tracking datasets of 2-branch and 5-branch risky choice (total 289 subjects). The best-fit parameters of our model provide strong evidence for interactive attention. These parameters also allow our model to predict key patterns in attention, choice, and response time. Although some participants do not display interactive attention, the ones that do are more likely to behave according to SEUT. Finally, we show that the proposed interactive attention mechanism has a variety of adaptive properties that make it useful for risky decision making. Models without interactive attention make absurd mistakes, not observed in human participants. By directly relating established theories in psychology and neuroscience to those in economics, our results open up new research directions at the interface of the social and natural sciences. Future work can test whether our mechanism is also at play in intertemporal and social choice, which also involve multiplicative utility maximization.

R-5: This time is different: On similarity and risk taking after experienced gains and losses

Steve Heinke[1], Adrian Leuenberger[1], Jöerg Rieskamp[1]
[1]University of Basel

How do experienced prior loss or gains affect risk-taking? A large literature reports significant but seemingly inconsistent effects of prior outcomes on risk-taking. We resolve these inconsistencies by proposing a similarity based mechanism determining which outcomes are jointly evaluated and state conditions under which we expect no behavioral changes. In line with our theory, we find in a pre-registered experiment, that the less similar a prior decision situation is in task-relevant dimensions, the weaker is the effect of the prior outcomes on the current decision; variation in non-task relevant dimensions will not change the impact of prior outcomes.

R-6: Association of the altered sense of time in the seconds-to-minute range with intertemporal choices across time-horizons

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

Delay-discounting studies in neuroscience, psychology, and economics have been mostly focused on concepts of self-control, reward evaluation, and discounting. Another important

relationship to consider is the link between intertemporal choice and time perception. We recently compared subjects' discount factors when faced with short (seconds) and long (days) time-horizon decisions. The short delays were experienced by the subjects after each choice. Only one long delay was experienced, after all the choices were made. While the subjects' behavior was correlated across the two tasks, there remained a substantial amount of unexplained variance. Here, following preregistration, we presented 50 college students with timing exercises on the range of seconds to minutes and intertemporal-choice tasks on both the time-scale of seconds and of days. We hypothesized that individual differences in time perception would influence decisions about short experienced delays but not long delays. While we found some evidence that individual differences in internal clock speed account for some unexplained variance between choices across time-horizons, overall our findings suggest a nominal contribution of the altered sense of time in intertemporal choice.

R-7: Present-biased time preference for monetary rewards - a meta-analysis

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

Background: 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. The model captures failures to stick to plans with actions that are costly now and have benefits in the future, such as saving, completing work, and exercising. Existing empirical evidence on quasi-hyperbolic discounting has largely relied on studies using monetary rewards to elicit the present bias parameter. At the current state of literature, there is no consensus that people are present biased towards monetary rewards. Objectives: (i) determine the average value of present-bias for monetary rewards in the existing studies, (ii) determine if there is a publication bias in the estimates, and (iii) explain the heterogeneity in findings across studies. Method: To deliver an unbiased meta-analysis, our identification and selection of papers are guided by unambiguously defined inclusion criteria. To create our dataset, we searched all the major research paper databases using the following methodology keywords: "elicit*", "estimate*", "experiment*", "measur*", "comput*", "test*", and the following topics keywords: beta-delta", "dynamic consistency", "dynamically consistent", "dynamic inconsistency", "dynamic inconsistent", "hyperbolic discount", "non-constant discount", "present bias", "future bias", "quasi-hyperbolic", "time consisten*", "time inconsisten*". Out of 2351 articles from this search, both published and unpublished, 716 survived a title and abstract screening with inclusion criteria: "relates to time preference", and "has empirical contents". Through full-text screening, 110 articles survived with inclusion criteria "reported point-estimated present bias parameter", and a further 78 survived with inclusion criteria "present bias can be inferred". We individually contact authors of the 78 articles for the access to their original data to estimate

present bias parameter and when possible the estimates will be included in our dataset for meta-analysis. Discussion: We calculate the "average" present bias for monetary rewards and investigate how types of study (e.g. experimental vs. observational; field vs. lab), methodology (e.g. convex time budget, longitudinal), subject pool, payment method, and publication status of a study affect the estimated magnitude of present bias.

R-8: The Effect of Investment Position on Belief Formation and Trading Behavior

Kevin Trutmann[1], Steve Heinke[1], Jörg Rieskamp[1]
[1]University of Basel

Understanding how prior outcomes affect the incorporation of new information into expectations is critical for the study of investment decisions: Three strands of literature separately document stronger belief updating in the case of prior losses, when being invested at all, and when the information indicates a price movement towards the direction of the initial buying price (i.e. mean reversion). We combine these three separate strands of literature and use a pre-registered experimental dynamic investment task. The 192 participants could invest in an asset whose price movements were autocorrelated and thus informative about future price developments. In each round they could not only decide whether or not to invest in this asset, but we further elicited participants expectations about the increase of the price to track the way they updated their beliefs throughout the task. In line with our reasoning we find that the belief formation of invested participants shows a pattern of reversal to the buying-price, which is more pronounced for loss positions. As investment decisions follow the expectations, an investment with gains is more likely to be liquidated than a losing one. In sum, we find evidence for a learning based explanation of the Disposition Effect (DE). In a second phase of the experiment, we endogenously manipulate the required amount of learning by providing additional information. Participants were either shown the objective probability of a price increase in the next round, the probability as calculated by a Bayesian updater or they repeated the no-information condition of the first phase. We find that provision of full information, i.e. the objective probability of a price increase, brings the investment decisions closer to the rational benchmark and the more learning is required, the stronger the deviation from the rational benchmark. These treatment effects underline the importance of understanding expectation formation in investment decisions if one wants to improve financial decision making.

R-9: Risky decision-making in the laboratory mediates the influence of amygdala activity on real-world financial behavior among individuals with low-SES backgrounds.

ranjita poudel[1], Michael Tobia[1], Michael Riedel[1], Taylor Salo[1], Jessica Flannery[1], Lauren Hill-Bowen[1], Angela Laird[1], Carlos Parra[1], Matthew Sutherland[1]
[1]Florida International University

Socioeconomic disparities are associated with poor health and economic outcomes including substance abuse, behavioral addictions, and greater consumer debt all potentially related to altered risky decision-making (DM) brain processes. However, the associations between risky-DM-related brain activity, laboratory-based risky behavior, and real-world financial outcomes remain poorly understood. As such, we employed a variant of a risky-DM paradigm (the Balloon Analogue Risk Task [BART]) and collected behavioral and fMRI data from individuals (n=27) with a low-SES background (income<$20,000/year) to interrogate risky-DM processes. We also tracked participants' real-life savings behavior while they participated in a 6-month savings program. We expected to observe individual differences in activation in reward-related brain regions during risky-DM and that such activity would correlate with a measure of BART performance: effective pump bet (EPB) and/or real-world asset building behavior (i.e., total savings). We further hypothesized that BART performance would mediate the relationship between brain activity and asset building behavior. Regarding BART-related brain activity, we observed increased activity during the DM task-phase in the bilateral amygdala, bilateral thalamus, bilateral striatum, and left frontal gyrus. Given the amygdala's role in emotional processing during risky-DM and its implications in psychopathology related to social adversity (including low-SES), we extracted brain activity estimates from three amygdala subregions (i.e., centromedial amygdala [CMA], superficial, and laterobasal) and performed correlations with EPB and total savings. We observed that right CMA activity was positively correlated with EPB (r[24] =0.44, p=0.03]. As EPB was also correlated with total savings (r[24]=0.58, p=0.003), we then conducted a mediation analysis to test if the influence of CMA activity (X) on savings (Y) was mediated by EPB (M). We observed that EPB fully mediated the influence of right CMA activity on savings. Specifically, when including EPB as a mediator, the CMA's direct effect on savings failed to reach significance (c' path: β=-0.18, p=0.4), whereas the indirect effect was significant (ab path: β=327.5, CI95%= 51.9, 799.7). These outcomes suggest that elevated CMA activity predicted increased ability to alternatively manage risks (i.e., EPB) and, in turn, this increased ability to alternatively manage risks predicted increased savings outcome. These findings suggest that individual differences in risky-DM laboratory performance and associated amygdala activity may be relevant factors for advancing real-world asset building behaviors.

R-10: The Effect of Delay Length on Time Preferences Across Adulthood

Shelby Leverett[1], Christopher Garza[1], Kendra Seaman[1]
[1]University of Texas at Dallas

Objective: The extent to which the larger, later option is discounted varies largely between people. Age has been hypothesized to explain some of this variance, but studies of age on

discounting are mixed. One reason for this heterogeneity could be the use of a mixture of time delays across studies, ranging from days to years. We posit that uncertainty about future longevity may lead to greater discounting in older adults for very long time delays (i.e. years). This study investigated how the delay length influences time preferences across adulthood. We hypothesize that age differences for time preferences will only appear with longer time delays. We do not expect to see any age differences in discounting with short time delays (i.e., days, weeks, months), but for longer time delays (i.e., years), we predict older adults to discount more than younger adults. Methods: Healthy participants (N = 288, Age: M=54.37, SD= 16.68, Range= 25-84) were recruited to complete a series of online surveys. Participants completed a demographic survey, the Future Time Perspective Scale, a savings & safety net questionnaire, and a delay discounting task. The task systematically varied the time interval between the smaller, smaller sooner and larger, later option. A combination of 12 time intervals (1 day, 4 days, 7 days, 1 week, 2 weeks, 4 weeks, 1 month, 6 months, 12 months, 1 year, 5 year, and 10 years) and 3 hypothetical discount rates (k= 0.1, k= 0.05, and k=0.005) were used to create 36 delay discounting trials. A hyperbolic discounting function was used to estimate an appropriate future monetary amount that corresponded to these time delays and discount rates. Discounting was measured as the proportion of smaller, sooner options chosen. Results: When delay was treated as a continuous predictor (number of days), there was a significant age x delay interaction on choice in the predicted direction. At the longest delay (days=3650), older adults discounted more than their younger counterparts. However, for every other delay length, older adults discounted less than younger adults. Contrary to our predictions though, greater discounting in older adults for the longest time delay was not mediated by their sense of uncertainty (or lack thereof) about their future. Conclusion: Our results suggest that older adults only discount when there are extremely long time delays (i.e. ~10 years). Future research could investigate time delays between 1 year (where age differences are not typically observed) and 10 years (where we observed age differences).

R-11: Neural Mechanisms of Reward Processing During Online Vickrey Auctions: An Event-Related Potential Study

Alice Newton-Fenner[1], John Tyson-Carr[1], Hannah Roberts[1], Jessica Henderson[1], Adam Byrne[1], Danielle Hewitt[1], Nick Fallon[1], Yiquan Gu[1], Olga Gorelkina[1], Yuxin Xie[2], Athanasios Pantelous[3], Timo Giesbrecht[4], Andrej Stancak[1]
[1]University of Liverpool, [2]Southwestern University of Finance and Economics, [3]Monash University, [4]Unilever

As the digital marketplace is expanding, it is becoming increasingly important to understand how people make purchasing decisions in online contexts. The online auction giant eBay uses a Vickrey auction format, where the highest bidder wins but the price paid is the second-highest bid. This type of auction gives bidders an incentive to bid their true subjective value, or willingness-to-pay, differentiating it from other auction paradigms. From previous studies,

reward outcome-related potentials (P200, FRN, P300, LPP) are known to be differentially sensitive to reward valence, magnitude, salience and context in gambling tasks. The present study examined, for the first time, the event-related potentials (ERPs) occurring during outcome processing in Vickrey auctions. Electroencephalographic feedback-related potentials were predicted to differentiate between three conditions: outbid (loss), win by a large margin (bargain) and win by a small margin (snatch). Twenty-eight healthy participants performed a Vickrey auction task, bidding on 300 household items grouped into low (£3-5) and high (£7-9) retail price. Participants bid against an anonymous computerised opponent, whose bids were modulated to be 10-30% (bargain), 70-90% (snatch), or >100% (loss) of participant's bids. Auction outcomes were associated with a feedback-related negativity (FRN) potential with a spatial maximum at vertex (251-271 ms). Bargain trials resulted in a greater FRN amplitude than snatch trials. The late positive potential (LPP) (361- 600 ms) also differentiated between bargain and snatch conditions, with bargain being more similar to the loss condition than snatch. The P200 and P300 components were sensitive to retail price, with greater amplitudes in response to higher value items. Bid outcomes in VA modulate a set of ERPs including the FRN and LPP. The FRN and LPP differentiated between wins / losses, and also between the two types of win. This suggests a sensitivity to instantaneous evaluations of relative value of gain. The P200 and P300 were not sensitive to the margin of win, but encoded the magnitude of the reward. Future research will unpack how the social competitive component of the Vickrey auction contributes to different ERP components.

R-12: The repulsion effect in preferential choice and its relation to perceptual choice

Mikhail Spektor[1], David Kellen[2], Karl Christoph Klauer[3]
[1]Universitat Pompeu Fabra, [2]Syracuse University, [3]University of Freiburg

When choosing among more than two options, humans and other animals frequently violate the independence principle of rational choice theory, according to which relative choice probabilities should not be affected by the composition of the choice set. The most prominent such violation is the attraction effect that occurs when a decoy that is dominated by one of the other options increases its choice probability. Recently, the reversal of the attraction effect, the repulsion effect, has been reported in perceptual choices. The study investigated if repulsion effects are restricted to perceptual choices, whether they are related to the attraction effect in preferential choice, and which cognitive processes are involved. The present study used eye tracking and a within-subject design of a perceptual task and its preferential counterpart. A total of 55 participants completed two 1.5-hour sessions at least one week apart with 700 decisions between three options in each session. The options were characterized by two attributes represented by horizontal and vertical bars. In the perceptual condition, individuals chose the option that had the highest sum of attributes. In the preferential condition, the attributes represented the outcomes of a binary lottery with equiprobable outcomes. The trials were designed such that the decoy was dominated by each of the other two options on half of

all trials, the standard setup for the attraction effect. In both conditions, the decoy affected the choices of the other two options. In line with recent research, we observed a repulsion effect in the perceptual condition, $t(54) = 5.73$, $p < .01$. However, we also observed the repulsion effect in the preferential condition, $t(54) = 2.37$, $p < .05$. The eye-movement patterns suggest that people integrate information within options more often than within attributes (across options). A logistic regression identified that the number of within-attribute comparisons is an important predictor of the attraction effect and that the lack of such comparisons is the main driver of the repulsion effect in our data. The multiattribute linear ballistic accumulator model, one of the most successful models of the attraction effect, was fitted to the data and was able to account for the data only slightly above chance level. Our study highlights the importance of perceptual task features in preferential choices, demonstrates the repulsion effect in a preferential-choice setting, relates it to the repulsion effect in the perceptual domain, and provides a process-level explanation for it. Furthermore, it corroborates the need for computational models that are able to account for repulsion effects.

R-13: Decision-Making Under Uncertainty in Borderline Personality Disorder Patients

Mathi Manavalan[1], Xin Song[1], Iris Vilares[1]
[1]University of Minnesota - Twin Cities

Existing research has developed a working understanding of borderline personality disorder (BPD) patient traits and behavior in everyday life, but the subtleties of their cognitive processes during decision-making remains unclear. To understand how reliance on previous experiences (priors) versus current sensory in¬formation (likelihoods) in the decision-making process may differ for those with BPD in comparison to those within neuro-typical population, we implemented a coin-catching behavioral task with varying levels of prior and likelihood uncertainty. This task was completed by eighteen BPD patients and twenty-three control participants. We hypothesized that, in accordance to typical BPD characteristics, BPD patients will rely significantly more on likelihood information even when likelihood information is more unreliable than prior information. Analyzing the results using Bayesian statistics, we found evidence suggesting that both the BPD patient group and the neuro-typical control group utilized prior and likelihood information similarly in decision-making under uncertainty. We theorize that BPD characteristics that are prominent in social interactions may not exactly replicate in non-social settings.

R-14: Decision-Making Under Uncertainty in Major Depression Patients

XIN SONG[1], Iris Vilares[1], Mathi Manavalan[1]
[1]University of Minnesota, Twin Cities

Background: Substantial evidence has suggested that major depression is associated with a dysregulated dopamine system, which plays a pivotal role in decision-making under uncertainty. In the Bayesian decision framework, people make decisions using information from both past beliefs (prior) and current sensory information (likelihood). Previous research has proposed that dopamine enhances the weight given to current sensory information (sensory weight) versus prior beliefs, yet how much this relationship holds true in depression remains a topic under debate. Study's objective: To examine whether depression patients have decreased sensory weight in decision making because of dopamine dysregulations. Method: We used a visual coin-catching task in which uncertainty in both prior and sensory information varied. 13 depression patients and 18 neuro-typical controls participated in this study, and decision-making strategies during the task were modeled by Bayesian statistics. Results: In line with our prediction, depression patients preserved the ability to learn both prior and sensory information uncertainty, comparable to healthy controls. In contrast to our prediction, depression patients did not show decreased reliance on sensory information compared to controls. Conclusion: Our results suggested that depression does not induce a universal alteration in decision-making strategies under uncertainty. Our study suggests that depression may not results in deficits in uncertainty processing at the sensory-motor level, despite its correlation with dopamine dysregulations.

R-16: Discovering the dynamic correlation between personality determinants and risk preferences in the field

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

A variety of personality determinants and emotional states have been linked to the individual differences of the risk preferences. However, most variables reported in the literature were investigated independently and tested without repeated measurements. In order to understand the interaction between these traits and states and risk preferences, we performed longitudinal examinations of emotional states and personality traits from the following domains: attitude toward uncertainty, impulsivity and the level of psychological distresses. All subjects were recruited in the US on CraigsList. We recruited a 50 subject "discovery cohort" and another 50 subjects as the "replication cohort". In the discovery cohort, our goal was to identify candidate correlations between the measured traits, states, and risk preferences. We employed a mobile smartphone-based experimental platform (Linkt) for gathering daily data from our participants. Each participant was tested with multiple instruments presented over a two-month window. Subjects completed instruments for about 5 minutes per day (two tasks and/or questionnaires). We gathered 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). Over the course of the study, we gathered a total of 13 personality inventories. We found that our data collection system yielded

a high level of compliance, 49 out of 53 subjects finished all instruments. Here we present data showing how the personality traits are correlated with willingness to take risks and its relationship with the risky choice task response times (RTs). We found that risky choice task RTs show a significant negative correlation across subjects with level of psychological distress measured from brief symptom inventory. The negative correlation between RTs and anxiety level is shown in both trait-like and state-like anxiety measured from the state-trait anxiety inventory. In addition, RTs is also negatively correlated with the level of pessimism in the life-orientation test and the willingness to take financial risk in the domain specific risk-taking scale. We found that the ratio of choosing a risky option over a safe option is positively correlated with the psychological distress as well as the willingness to take financial risk. Our findings demonstrate that our data collection system is effective in capturing the relationship between the variables. For the next step, we will be testing these effects with our replication cohort. Disclosure: PWG is an officer and stockholder in Linkt Datacubed Health.

R-17: The effects of dopaminergic psychostimulants on complex problem solving

Elizabeth Bowman[1], David Coghill[1], Carsten Murawski[1], Peter Bossaerts[1]
[1]The University of Melbourne

The role of dopaminergic systems, particularly the mesocorticolimbic dopaminergic pathways and their associated cortical projections, in decision-making and goal-directed behaviours has been of increased research interest in recent years. Medications that target these pathways are often used in disorders of attention and cognition, and are now also often diverted to non-medical, off-label, and illicit uses in the hope that they will enhance cognition and focus in healthy people. However, research to date reveals a mixed picture of efficacy in this aspect, particularly with regard to performance in higher-cognitive tasks. We completed a repeated-measures, double-blinded, placebo-controlled single dose trial (PECO: ACTRN12617001544369, U1111-1204-3404) to examine the actions of three indirect dopamine agonists on complex problem-solving performance of healthy adults. Forty participants received 15 mg dextroamphetamine, 30 mg methylphenidate, 200 mg modafinil, or placebo, counterbalanced across 4 testing sessions. Testing sessions were held at least one week apart. In each session, participants completed 16 trials of the Knapsack Task (eight unique instances presented twice each), and the CANTAB Simple and Five-choice Reaction Time, Spatial Working Memory, Stockings of Cambridge, and Stop Signal Tasks. The Knapsack Task is a NP-hard combinatorial optimisation task, requiring the participant to explore different combinations of items to discover an optimal (maximal) value solution under a specified weight and time constraint. Mixed-effects modelling of the knapsack decision sequence data revealed that in general, participants took longer to submit solutions in the active drug conditions yet were less likely to find the optimal solution. Participants who performed below-average with placebo explored more solutions while in the active drug conditions. Yet their performance still decreased. Above-average participants reduced performance too, as they actually explored less. Closer

inspection revealed that the drugs decreased productivity significantly and economically; productivity is the increase in value of the knapsack per item move per unit time. Initial choices, a good predictor of eventual performance in a trial, became more random under active drug conditions. These findings suggest that dopaminergic stimulant medications may increase healthy adults' motivation to spend more time and explore more combinations in pursuit of the solution to a computationally complex problem. However, this does not always result in enhanced performance because exploration becomes more random, which decreases performance.

R-18: The Influence of Hormones and Personal Traits on the Propensity for Risk-Taking

Ursa Ferjancic[1], Aljosa Valentincic[1]
[1]School of Economics and Business, University of Ljubljana

In the present study, we conducted an experiment to test the relations of the propensity for risk-taking, personal traits (i.e., optimism, personality, and lifestyle characteristics) and hormones (i.e., testosterone and cortisol) using a sample of 36 post-graduate students (21 women, 15 men; Mage=23.6, SDage=1.30). We measured risk propensity with a self-reported (Domain-specific Risk-Taking Scale - DOSPERT), and behavioral measures (Balloon Analogue Risk Task, Eckel & Grossman Risk Task, Holt & Laury Measure of Risk Aversion). Laboratory assistant from a certified laboratory collected saliva samples and analyzed the concentrations of testosterone and cortisol in the saliva. We used the digit ratio (2D:4D) as a proxy for prenatal testosterone exposure and facial Width-to-Height ratio (fWHR) as a proxy for investigating the relations among facial characteristics and the propensity for risk-taking. We measured personality characteristics using the Zuckerman-Kuhlman Personality Questionnaire, while we obtained optimism and lifestyle characteristics by a demographic questionnaire. We applied the Sensation-seeking scale to measure the sensation-seeking trait. The Committee of Ethics and Research at the School of Economics and Business of the University of Ljubljana approved the experimental design and all the related procedures. Before participation in the experiment, all subjects agreed to participate by signing a written informed consent to participate. Our results show support for the dual-hormone hypothesis, but only when we employed the Holt & Laury Measure of Risk Aversion as a dependent variable. Additionally, we find testosterone and cortisol levels are positively related to risk propensity in the financial domain (DOSPERT). We find that leading traits of general personality patterns for risky behavior are impulsive sensation-seeking and neuroticism-anxiety, which are positively and negatively correlated to distinct risk propensity measures, respectively. Also, the results show many positive relations between the sensation-seeking traits and various risk-propensity constructs. Finally, we find positive associations for optimism, alcohol consumption, and smoking for some risk-taking constructs as well. Our results do not exhibit any statistically significant correlations of fWHR and 2D:4D to the risk propensity measures, respectively.

R-19: The role of individual differences in information processing on risky decision-making

Kevin da Silva Castanheira[1], Ross Otto[1]
[1]McGill University

Risky decision-making reflects time-dependent information processing where actions' consequences and their associated probabilities of occurring are compared. While previous work has shown that limiting deliberation time increases the extent to which people are sensitive to outcomes' framing (i.e. losses vs gains) less is known about the effect of encouraging longer deliberation times. Here, we sought to better understand how both outcomes and their associated probabilities inform decisions in a time-dependent manner by manipulating the amount of time for information processing. Additionally, given the considerable variance in how quickly individuals process information, we also sought to test whether individual differences in information processing shape risk preferences. In two experiments, participants (N=100) completed a risky-decision making task with problems framed either in terms of gains or losses under two conditions of available time for deliberation. Responses and response-times were fit to a Drift Diffusion Model. Results provide evidence for a starting bias for responses in accordance with the framing effect: suggesting that participants primarily use outcome information to guide their choices. In a second experiment, participants were asked to use all the time available for deliberation. In line with the model's predictions, data in our second experiment show that both longer deliberation times and faster information processing were associated with a weaker effect of framing and an increased effect of outcome probability on choice, in a manner consistent with the Fourfold pattern. Together, the results suggest that both deliberation time and processing speed limit the amount of information considered in risky choices and offer a novel framework for understanding the dynamic nature of risk preferences and how information processing shapes choice.

R-20: Pupil dilation reflects individual differences in temporal expectations during persistence for delayed rewards

Karolina Lempert[1], Sakshi Sehgal[1], Joseph Kable[1]
[1]University of Pennsylvania

Objective: Individual differences in how long people persist in waiting for delayed rewards with uncertain timing have been linked to life outcomes such as academic success. Those individual differences have been attributed to variability in willpower, but they could instead be driven by differences in expectations about when delayed rewards will arrive. Previous work has shown that temporal expectations affect willingness-to-wait, but do temporal expectations underlie individual differences in willingness-to-wait? Here we address this question, using pupil dilation as a latent marker of expectations. Methods: Eighty-three subjects did a willingness-to-wait

task while eye tracking data were recorded. On each trial, they waited through a random delay (0.2-40 s) for a 10¢ reward. They did not know how long the 10¢ would take to arrive, and they could quit waiting at any time and move on to a new trial. Their goal was to maximize earnings in 20 min. Reward delays were distributed such that there was no optimal time at which to give up waiting. This allowed us to isolate idiosyncratic willingness-to-wait. Pupil analyses focused on two epochs in the trial: the 4 s right after subjects either got a reward or quit, and the last 4 s of the trial leading up to reward receipt or quitting. Results: Pupil dilation increased following rewarded, compared to quit, trials. This reward-related phasic pupil response was modulated by how long the subject waited for the reward. Specifically, subjects showed the largest physiological "surprise" response when the 10¢ arrived after a very short time or after a very long time. The wait time after which people were least surprised to be rewarded (i.e., minimum phasic pupil response) was associated with their willingness-to-wait: people who tended to wait longer for rewards were less surprised when rewards took longer to arrive ($r=0.38$; $p<0.001$). Pupil diameter then decreased through the waiting period of the following trial, with the decrease becoming gradually less steep: in the last 4 s of each trial, the slope of the pupil diameter was more positive on longer trials ($\beta=0.0003$; $p<0.001$). This relative increase in pupil dilation is consistent with the pupil's role in signaling decision uncertainty, which increases as time elapses. This relative increase was steeper in subjects prone to quitting early ($r=-0.38$; $p<0.001$), as if their uncertainty about waiting for the reward increased earlier. Conclusions: We quantified individual differences in willingness-to-wait for delayed rewards with unknown timing. Our pupillometry results suggest that those individual differences arise from differences in temporal expectations.

R-21: A 7-week dietary intervention increases delay discounting in an intertemporal choice task

Leonie Koban[1], Daniela Schelski[2], Marie-Christine Simon[2], Bernd Weber[2], Hilke Plassmann[1]
[1]INSEAD & ICM, [2]Universitaet Bonn

Objective: Growing evidence suggests important links between the gut microbiome--the bacteria and microorganisms that live in the gut--and host behavior and health. Yet, whether and how human decision-making can be influenced by changes in gut microbiome composition is unknown. Here, we tested whether delay discounting--how much people discount delayed compared to immediate monetary rewards--is affected by a dietary intervention targeted at altering the composition of the gut microbiome. Methods: We conducted a double-blind placebo controlled dietary intervention study. Half of the N=117 male adult participants were randomly assigned to daily intake of a synbiotic dietary supplement (combination of pre- and probiotics) for seven weeks. Participants in the placebo group received a non-active supplement. We used an incentive-compatible intertemporal choice task (ITC) to assess delay discounting both before (T1) and after (T2) the intervention. Delay discounting was quantified as the percentage of smaller sooner choices (%SS) and by fitting the parameter k to estimate

individual hyperbolic discount functions. Results: We found a significant group difference in the percentage of smaller sooner (%SS) choices (p = 0.0071) after relative to before the intervention. Planned comparisons revealed that, while no change was observed in the placebo condition, participants in the intervention group showed a significant increase in %SS choices following the intervention (p = 0.0015). Parallel results were found for the fitted hyperbolic discount parameter k (log-transformed), with a significant group difference in the T2-T1 change in log(k) (p = 0.034), and a significant increase in log(k) in the intervention (p = 0.0041), but not in the placebo group. Conclusions: Our study showed that a seven-week intake of pre- and probiotics--thought to alter the composition of the gut microbiome--can lead to increased discounting of delayed rewards, i.e. more impatient choices. These results reveal a surprising link between dietary factors and impatience. The effects of dietary changes on discounting may be mediated by changes in dopamine and serotonin precursors and brain systems linked to these neurotransmitters, which are known to modulate decision-making. Future studies should further investigate the mechanisms of these effects and whether interventions targeted at the microbiome can improve clinical conditions that are characterized by abnormal discounting.

R-22: Effects of cognitive load on initial and sustained uncertainty-driven exploration in a multi-armed bandit task

Vanessa Brown[1], Michael Hallquist[2], Michael Frank[3], Alexandre Dombrovski[1]
[1]University of Pittsburgh, [2]University of North Carolina, [3]Brown University

Introduction: How humans resolve the explore/exploit dilemma is unclear, especially under conditions of uncertainty and cognitive load that mirror real-world demands. To investigate if humans are uncertainty-seeking or -averse, and if cognitive load modulates this tendency, we manipulated two forms of cognitive load (whether participants needed to track past outcomes and the number of options available) to understand how these factors affected initial exploration and the shift from exploration to exploitation during learning. Methods: 95 participants completed a multi-armed bandit task with 8 blocks of 30 free choice trials each. Blocks differed in the number of options (4 vs. 8) and whether the outcomes (win/no win) of previous choices were shown. To manipulate initial uncertainty, each block started with a set of forced choice trials. In some blocks each option was sampled once during forced choice trials, while in other blocks options were sampled unevenly. Analyses of initial exploration examined whether participants' initial free choices were exploratory (option not chosen during forced choice trials), exploitative (option always rewarded during forced choice trials), or random (option not consistently rewarded during forced choice trials). Analyses of the transition to exploitation with learning examined how choices narrowed over time. Results: Participants' initial exploration was significantly altered by cognitive load: with a greater number of options, participants were less exploratory and more random (exploitation unaffected; $\chi 2=11.1$, p=0.03), while if previous outcomes were shown, participants were more exploratory and less exploitative (random choices unaffected; $\chi 2=25.0$, p < 0.001). During learning, participants

narrowed choices over time (t = 7.97, p < .001), showing a transition from exploration to exploitation. This effect was more pronounced with a greater number of options (t = 8.78, p < .001), but less in blocks with uneven initial sampling (t = -6.50, p < .001). Additionally, showing previous outcomes led to more rapid narrowing with a smaller number of options (t = 3.20, p = .001) but less narrowing over time with more options (t = -4.81, p < .001). Conclusions: Participants' uncertainty attitudes differed based on cognitive load. Different costs also affected exploration differently: the increased cognitive cost when learning about many options caused inefficient exploration early in learning and over-exploitation of certain options later, possibly resulting from poor maintenance of values, while needing to maintain outcome values in memory lead to an emphasis on exploitation at the expense of fully exploring options.

R-23: A potential neural signature of impaired self-control pertinent to excessive trading: Low functional connectivity between dorsolateral and ventromedial prefrontal cortex

John Haracz[1]
[1]Indiana University

Objective: Neuroeconomic studies have begun to examine heuristics (i.e., simple rules that guide behavior while minimizing effort) and biases (i.e., systematic behavioral deviations from rational-choice models [Mobbs et al., 2018]). Haracz (2020) proposed that inadequate deliberation could bias traders toward simply using 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. Ogawa et al. (2014) used functional magnetic resonance imaging (fMRI) to study neural mechanisms of a bias to buy stocks (i.e., a bubble bias) while subjects traded in lab markets with stock-price bubbles. In the studied lab-market conditions, a bubble bias was found only during the bubble period of exposure to historical Lehman Brothers stock-price changes from 2003-2008. However, Ogawa et al. (2014) failed to discern evidence from their data suggesting that the bubble bias was associated with a neural signature of impaired self-control. Therefore, their findings are presently re-examined. Methods: In Figs. 2C and 2D of Ogawa et al. (2014), fMRI functional-connectivity data were examined to determine whether a bubble bias was associated with the previously reported neural signature of decreased self-control (Hare et al., 2014; Maier et al., 2015). Results: Among all lab-market conditions (Ogawa et al., 2014), low dorsolateral prefrontal cortex (DLPFC)-ventromedial prefrontal cortex (VMPFC) connectivity was specific to the bubble period of Lehman Brothers stock prices. Among all conditions, the finding of a bubble bias was specific to the bubble period of Lehman Brothers stock prices. Conclusions: The low DLPFC-VMPFC connectivity associated with a bubble bias (Ogawa et al., 2014) resembles findings of low DLPFC-VMPFC connectivity linked to impaired self-control when subjects were exposed to a food reward (Maier et al., 2015). Low DLPFC-VMPFC connectivity also was found when subjects failed to resist readily available monetary rewards in an intertemporal choice task (Hare et al., 2014). Therefore,

subjects with a bubble bias show connectivity changes that may represent the neural signature of impaired self-control. Neuroimaging studies should determine whether less-profitable investors who trade stocks excessively (Barber and Odean, 2000, 2001, 2002) may show similarly decreased DLPFC-VMPFC connectivity. These results encourage further research on neural mechanisms underlying the bubble bias and traders' self-control problems that may promote asset-price bubbles.

R-24: Investigating the Neural Correlates of the Risk Preference and Information Acquisition in the Description-Experience Gap

Elizabeth Beard[1], Jason Chein[1], Vinod Venkatraman[1]
[1]Temple University

Objective: The Description-Experience (D-E) Gap demonstrates that risk preferences can change as a function of whether the same information is described or learned from experience. Critically, while human studies often use descriptive tasks to study risk preferences, non-human primate studies use experiential learning. The current study aims to address the impact of these methodological differences by studying neural processes associated with decisions from description (DFD) versus experience (DFE), and the resultant D-E Gap. Methods: 24 young adults completed a novel, within-subjects paradigm while undergoing fMRI. To compare differences across decisions from DFD and DFE, participants completed 36 choices in each format, with trials divided into three separate phases. In the acquisition phase (15-19s), participants either sampled outcomes from two separate decks repeatedly (DFE) or viewed visual information about the probabilities and outcomes of the two alternatives (DFD). To balance memory requirements across formats, in DFD trials each probability and outcome from the two alternatives was presented sequentially in a randomized order. During the think phase (3s) participants were instructed to consider which of the two options they would choose based on the information acquired. In the last phase of the trial, participants made a selection between the two available options (2s max). Results: Consistent with previous research, we found greater risk taking in DFE than DFD. Participants were more likely to underweight rare unfavorable outcomes in DFE and overweight rare favorable outcomes in DFD. In DFE, sampling behavior did not have an effect on choice. Neither the total number of samples, nor the number of switches (switching from the sampling of one option to the other) was significantly related to selecting the risky option. Preliminary fMRI analyses reveal stronger activation in neural regions associated with episodic and working memory when acquiring information for DFD, indicating heavier memory demands in assimilating information about description. However, during the think phase, we found greater activation in the ventral striatum for DFE than DFD, consistent with value integration and increased risk preference in DFE. Discussion: Our results suggest that working and long-term memory may be important factors in explaining the D-E gap. Forthcoming analyses will examine the relationship between neural processes of information acquisition with information integration during think phase. These findings will also

provide critical insights into neural differences between human and nonhuman primate on risk preferences.

R-25: Biased Sequential Sampling underlies the Effects of Time Pressure and Delay in Intertemporal Choice

Fadong Chen[1], Jiehui Zheng[1], Lei Wang[1], Ian Krajbich[2]
[1]Zhejiang University, [2]The Ohio State University

Intertemporal choices - decisions with consequences that play out over time - pervade our everyday life. Such choices range from the mundane to life-changing decisions about health and saving behaviors. Dual-process theories of intertemporal choice assume that decisions result from competing dispositional preferences that evolve over time. However, debates about whether impatience (self-control failure) derives from intuition or controlled reflection have preoccupied economists and psychologists for decades. Here, we argue that intuition can be seen as biases in a sequential sampling process underlying intertemporal decision making, and we investigate how these biases influence choices and explain individual differences in the effects of time pressure and delay. In the experiment, participants (N=126) made 300 binary decisions about how to decide between a sooner but smaller reward (SS) and a later but larger reward (LL) in time-free (no time limit), time-pressure (2 s limit), and time delay (10 s wait) conditions. Fitting a hyperbolic discounting model to the choice data revealed that participants were heterogeneous in whether they were predisposed to money or time. In particular, compared to the time-free condition, patient participants (N=94, classified by the critical discount rate of the game) become more patient under time pressure (two-sided Wilcoxon signed rank test, $p<0.001$) and less patient under time delay ($p=0.003$), while impatient participants (N=32) did the opposite (time pressure: $p=0.091$; time delay: $p=0.031$). The drift diffusion model (DDM) starting-point and intercept biases in the time-free condition explained the change in preference across time pressure and delay conditions (two-sided Spearman correlation tests, starting point $rho=0.294$, $p=0.004$; intercept: $rho=0.381$, $p<0.001$). We validated our results using out-of-sample predictions, revealing that the biased DDM has higher predictive power (Cramer's lambda) than the unbiased DDM (two-sided Wilcoxon signed rank test, $p <0.001$). These findings help to explain inconsistencies in our understanding of the cognitive processes of intertemporal decision making.

R-26: Dynamic Representation of the Subjective Value of Information

Kenji Kobayashi[1], Sangil Lee[1], Alexandre Filipowicz[1], Kara McGaughey[1], Joseph Kable[1], Matthew Nassar[2]
[1]University of Pennsylvania, [2]Brown University

Adaptive information seeking is critical for decision making. Economic theory suggests that agents should seek information based on its instrumental benefit, i.e., to the extent it improves future actions and helps utility maximization. This prediction has been supported by behavioral evidence that human information-seeking behavior is primarily driven by the instrumental benefit. However, neural processes underlying valuation of information based on instrumental benefit is little understood, particularly when its instrumental benefit changes dynamically over time. Here we show that subjective value of information (SVOI) signals in human brains are adaptively updated to track changes in instrumental benefit. In an fMRI experiment, human participants (n = 15) performed a variant of the classic beads task. They were presented with a variable number of beads that were drawn from one of two jars. Each jar contained a 60/40 mixture of visually distinct bead types, but the jars differed in which bead type was in the majority. Participants were asked to make a bet on the majority bead type in the jar from which the beads were drawn. We introduced reward asymmetry, such that a correct bet on one jar type (manipulated across blocks) was rewarded with larger monetary reward than the other jar type. Prior to the bet, participants were allowed to purchase information and draw more beads from the jar. Theoretically, the instrumental benefit of the information for utility maximization is highest when slightly more low-reward beads have been drawn. In order to elucidate the neural processes of SVOI updating, we presented participants with one additional bead after the initial presentation of beads but before the information purchase phase. Participants' information purchases were biased in the theoretically predicted direction, such that they drew beads more often when slightly more low-reward beads had been drawn. This suggests that SVOI is primarily determined by instrumental benefit as normatively prescribed. At the neural level, upon the initial presentation of beads, SVOI was represented in preSMA, right anterior insula, and right middle frontal gyrus (MFG). SVOI representation in right MFG is consistent with a previous report (Kobayashi & Hsu, 2019), providing convergent evidence across different tasks. Lastly, these regions not only represented the initial SVOI but also responded to its updating; BOLD signals in these regions were increased when the additional bead updated the SVOI positively, and decreased when SVOI was updated negatively. These results demonstrate that neural SVOI signals track the instrumental benefit of information over time.

R-27: The surprising effect of social distancing on our perception: Coping with uncertainty

Giulia Piccillo[1], Job Van den Hurk[2]
[1]Maastricht University, [2]SCANNEXUS

Since the beginning of the COVID-19 crisis, our perception of the world significantly changed. In this paper we show the results of 3 studies that collectively illustrate a novel mechanism through which this has happened. We document the effect of social distancing on our perceptions, through the moderating effect of ambiguity aversion. In experiment 1 (146 subjects) we show that ambiguity aversion predicts illusory pattern perception, defined as identifying faces in white noise pictures. In experiment 2 (154 subjects) we show that ambiguity

aversion also predicts higher cognitive level illusory pattern perception, defined as belief in conspiracy theories. Experiment 3 (347 subjects) shows, through two uniquely timed questionnaires, that ambiguity aversion increases significantly from before to after the lockdown (due to the COVID-19 pandemic) for a sample of over 300 subjects. Finally, this difference in ambiguity aversion is no longer significant when we control for the drop in regular social contact over this period.

R-28: Decision-Making Among Individuals Reporting a COVID-19 Infection: Time, Probability, and Ambiguity Preferences

Alexandra Mellis[1], Candace Raio[1], Oluwatobi Olufeko[1], Angelica Vina-Albarracin[1], Paul Glimcher[1]
[1]NYU Langone Health

Overview: The COVID-19 pandemic imposes a universal threat to individuals' well-being, with individual risks of infection varying based on medical, social, and economic determinants of health. One question that has yet to examined is whether individuals' decision-making may also determine, or be influenced by, infection risk. Based on the previous literature linking economic choice patterns to health-related and clinical outcomes, here we directly test whether preferences for risk and ambiguity, as well as impulsive choice, differ in individuals as a function of COVID-19 infection. Methods: Between April 8 and July 22, 2020, a total of 232 participants consented to participate in "Our Covid Story" (www.datacubed.com/ourcovidstory). Participants were recruited via Facebook advertising and e-mail listservs, focused on the New York City area. Participation involved downloading a mobile application for administration of tasks and surveys, and completing 5-10 minutes of a rotating battery of assessments. These assessments included a complete demographic battery, self-report instruments, and gamified decision-making tasks, which included the 27-item Monetary Choice Questionnaire to assess temporal discount rate, and an abbreviated (35-item) version of the Levy task which distinguishes between preferences in risk and ambiguity. Results: Overall, 6% of individuals reported a likely or confirmed case of COVID-19. To assess decision-making, we calculated the proportion of immediate choices in the Monetary Choice Questionnaire to determine degree of impulsive choice; similarly, we calculated percent of lottery choices in the risk and ambiguity trials of the Levy task. In cross-sectional analyses, we observed that individuals who reported a positive history of COVID-19 made fewer risky lottery choices ($t=-2.81$, $p = 0.01$) and made fewer ambiguous lottery choices ($t = -2.22$, $p = 0.04$), suggesting lower risk and ambiguity tolerance in participants reporting COVID-19 infections. We additionally observed a small difference in the proportion of impulsive choices between these groups, with those reporting a history of COVID infection showing higher levels of impulsivity, although this group difference did not reach statistical significance ($t=1.23$, $p = 0.26$) in our initial sample. Discussion: Our longitudinal data collection is ongoing, to determine whether differences in decision-making between those who report a history of COVID and those who do not are cause or consequence of the disease. Our preliminary data support the application of econometric approaches to

understanding decision-making in COVID. Disclosure: PWG is an officer and stockholder in Datacubed Health

R-29: Neural correlates of reaction time variability in the human brain

Ashwin Ramayya[1], Vivek Buch[1], Andrew Richardson[1], Timothy Lucas[1], Joshua Gold[1]
[1]*University of Pennsylvania*

Introduction: Reaction-time (RT) variability is a fundamental property of decision-making, but the underlying mechanisms in the human brain are not fully understood. Using intracranial EEG (iEEG), we tested the hypothesis that RT variability in a simple sensory-motor task is reflected in trial-by-trial fluctuations of distributed, motor-related ramping signals, as predicted by accumulate-to-bound decision models.
 Methods: We recorded iEEG data from 23 patients with medically-refractory epilepsy while they performed a stimulus-detection task with long (1500 ms) or short (500 ms) delays before stimulus onset. We extracted high-frequency activity (HFA, 70-200 Hz) as a measure of local population firing rates from 2,709 widely distributed electrodes throughout cortex and the medial temporal lobe. We identified electrodes with motor-related signals based on temporal relationships between local peaks in neural activity and button-press RTs. We related features of these motor-related neural activity patterns, including pre-stimulus baseline activity, pre-response buildup rate and peak activity to RTs using multi-variate regression at each of these electrodes. We tested the significance of the resulting regression coefficients by averaging the associated t-statistics within each subject and studying their distribution across subjects.
 Results: Subjects showed faster RTs during long-delay trials compared to short-delay trials, with substantial trial-to-trial RT variability for each condition ($t(22) = 5.57$, $p < 0.001$; mean +/- s.d. = 419.3 ms +/- 55.0 vs. 457.02 +/- 61.0). Across all 23 subjects, we identified 216/2709 electrodes from multiple brain areas with motor-related activity that closely predicted RTs on a trial-by-trial basis (8.7%; χ^2 statistic = 16727, $p < 0.001$; n = 21/23 subjects). Across these electrodes, we found RT-related neural fluctuations in several neural features. Specifically, faster RTs were associated with higher pre-stimulus baseline activity ($t(20) = 7.81$), faster pre-response buildup rates ($t(20) = 8.27$), and lower pre-response peak activity ($t(20) = 8.14$, all p-values < 0.001, False Discovery Rate corrected). We found similar RT-related changes in neural activity when considering only short or long delay trials.
 Significance: Our finding that certain distributed, motor-related ramping signals in the human brain co-vary with RT is generally consistent with simple accumulate-to-bound decision models. Our findings of additional, within-condition RT-related fluctuations in baseline and peak activity suggest these processes include additional degrees of freedom that can contribute to stochastic variability in RT.

R-30: The Evolution of Self-Control in the Brain

David Jimenez-Gomez[1]
[1]University of Alicante

Self-control is the ability to control one's thoughts, emotions, and behavior. People who have more self-control are better educated, wealthier and healthier. If self-control is so beneficial, and given that humans are the result of an evolutionary process, this begs the question: why have we not evolved a perfect self-control? In this paper I argue that there is actually an optimal amount of self-control from the point of view of genetic fitness, and that is why humans have evolved a limited self-control. I present a theoretical framework that shows how limited self-control could have evolved as a mechanism to make humans behave against their own self-interest when such behavior would increase genetic fitness. I analyze the evolution of self-control in a principal-agent framework with two agents, System 1 and System 2, that represent the automatic and cognitive processes within the human mind, respectively. Based on the relevant evidence from neuroscience, I assume that System 2 has access to private information, but its utility cannot depend on all the relevant information. The principal can achieve the asymptotically optimal outcome by biasing the utility of System 2 (from which an endogenous conflict emerges) and simultaneously endowing it with a limited amount of self-control. The model explains several empirical properties of self-control (observed in experiments). The model is consistent with a "muscle model" of self-control, in which it is depleted in the short run but can grow in the long run. However, the model shows that glucose (or some other source of energy) would not need to be consumed to exert self-control, attempting to shed some light on the recent controversy in the "ego depletion" literature. The model also has welfare implications: humans have less self-control than what would be optimal for their well-being (according to their own evolved utility function), and this is exacerbated by evolutionary mismatch, due to the growth in tempting opportunities in our current society.

R-31: Hyperbolic Discounting Is Not Lack of Self-Control

David Jimenez-Gomez[1]
[1]University of Alicante

Hyperbolic discounting and lack of self-control are different phenomena that share the property of dynamic inconsistency. Despite being different, these two concepts are often not adequately differentiated, and have been used interchangeably as shorthand for dynamically inconsistent behavior. Both concepts have been used to explain low levels of saving, procrastination, demand for commitment, etc. I review the definitions and properties of both concepts, the neuroscientific evidence behind them, and their evolutionary origins. Based on these, I argue that hyperbolic discounting is a particular type of dynamically inconsistent time preference that affects all goods, is context-independent, derived from a single valuation system, and evolved in alignment with genetic fitness. In contrast, lack of self-control affects particular types of goods, is context-dependent, derived from multiple systems, and evolved

from an endogenous conflict between the individual's utility and genetic fitness. From these characteristics, I discuss the importance of properly differentiating these two concepts.

R-32: Adaptive self-control: Environmental volatility influences precommitment decisions

Lauren Sussman[1], Joesph McGuire[1]
[1]Boston University

OBJECTIVE: To avoid failures of self-control, people can precommit to a future course of action that agrees with their current preferences. By eliminating goal-incongruent alternatives, precommitment can reduce the need to exercise effortful and potentially unreliable self-discipline at the time of choice. While precommitment can be a beneficial self-control strategy, it is only useful when external contingencies are seen as relatively stable. Precommitting to a dynamic option in a volatile environment could result in a disadvantageous outcome. Here we tested the hypothesis that precommitment decisions are made by weighing the anticipated unreliability of one's own actions against the instability of the external environment. METHODS: Participants (N=64) performed a computerized task in which they repeatedly tried to select the vendor that would give them the highest price for a used book. They could either commit to a vendor on the basis of preliminary prices, or wait to make a last-minute choice on the basis of final prices. If they chose to commit, there was a risk the preference ordering of the two vendors would reverse (representing environmental volatility), and the risk was greater if the preliminary prices were more similar. If they chose to wait, their last-minute choice required a speeded response with only a 75% chance of producing the desired effect (representing the risk of a lapse at decision time). The optimal strategy was therefore to commit if the preliminary price difference was large and wait if it was small. RESULTS: As hypothesized, participants more frequently chose to precommit when the preliminary price difference was larger, which implied a stable preference ordering over the alternatives, and more frequently chose to wait when the preliminary price difference was smaller. This tendency developed over time, consistent with incremental learning of the task's structure. Participants also showed an overall bias to precommit more often than ideal. In addition, participants unexpectedly chose to precommit significantly more often when the average preliminary prices were higher. Task behavior was best fit by a mixed effects logistic model with significant main effects of initial option value difference, average value of the options, and trial number (as a metric of learning over time), as well as interactions between both trial and option difference and trial and average value. CONCLUSIONS: The results indicate that people can make precommitment decisions by adaptively evaluating a tradeoff between the volatility of the external environment and the anticipated unreliability of their own future actions.

R-33: Methods to Analyze Human Delay Discounting Data Translate Well to Rodent Data

Katherine Garland[1], Deborah Sevigny-Resetco[1], Suzanne Mitchell[1]
[1]Oregon Health & Science University

Delay discounting, or intertemporal choice, is the relative preference for smaller, immediate rewards over larger, later rewards. Delay discounting has been used as a measure of impulsivity, and has been linked to the various psychopathologies including addiction, ADHD and mood disorders. Identifying the neural circuits associated with this economic, cost-benefit tradeoff behavior is critical to understanding the linkage to psychopathology. Research addressing this issue has been conducted in humans as well as animals. If findings from animal studies are to be generalized to humans, it is imperative to verify that the characteristics of delay discounting are translatable between humans and animal models. Through a larger project identifying differentiating genes and RNA transcriptome networks associated with delay discounting, we have assessed delay discounting behavior in 300 genetically-heterogeneous adult rats (150 males and 150 females) by assessing choice between a small, immediate sucrose solution reward and 150 ul delivered following a delay of 0, 2, 4, 8, 16, or 24 seconds using an adjusting amount procedure similar to that used with human subjects. Adopting strategies developed from human studies, data from subjects demonstrating extreme choices and nonsystematic discounting were removed from the pool of data prior to analysis. The percentage of nonsystematic data was similar to past human studies. Analyses of delay discounting in rodents have historically focused on simple delay preference indices (main effects of delay in ANOVAs). In human studies, choices are described by fitting mathematical functions and fit statistics used to evaluate function efficacy. When our cleaned data was analyzed, it was found that, in aggregate, the rodents displayed a clear hyperbolic discounting curve, and this was also true when functions were fit for individual subjects. Fit statistics were excellent and compared well those from human data. These results support the contention that studies using rodents can provide comparable behavioral outcomes, supporting the idea that translational, neuroeconomic findings can be generated from rodent models of delay discounting.

R-34: Perceived versus Actual Virus Transmission Risk During the COVID-19 Pandemic

Alyssa Sinclair[1], Shabnam Hakimi[1], Matthew Stanley[1], Gregory Samanez-Larkin[1]
[1]Duke University

During the COVID-19 pandemic, individual decision makers have been challenged with weighing evolving information about risk that can have critical consequences for personal and public health. We sought to characterize risk perception related to COVID-19 in a nationally-representative sample of U.S. participants (N = 303). Participants completed surveys about risk perception and pandemic-related life disruption, as well as the Social Value Orientation (SVO) task, which categorizes participants as Altruistic/Competitive/Individualistic depending on choices in a point allocation game. To compare subjective risk perception with actual risk levels

in each participant's county at the time of the study, we extracted objective measures of exposure risk based on data from the Johns Hopkins Coronavirus Database and related these measures to self-reported preferences and behavior. First, we found a striking disconnect between perceived and actual risk. Actual risk in a participant's location had no association with either perceived risk ($R = 0.03$, $p = .62$) or willingness to engage in risky behaviors ($R = -0.01$, $p = .85$). Second, we found that the relationship between perceived and actual risk differed depending on SVO type ($F_{2,227} = 5.1$, $p = .007$): when actual risk was high, Altruistic participants tended to perceive greater risk while Individualistic and Competitive participants tended to perceive lower risk. Third, we also found a significant interaction between age and actual risk predicting perceived risk, such that older adults (age 60+) reported more accurate risk perception ($F_{2,233} = 4.97$, $p = .027$). Lastly, we found that participants who perceived less risk tended to be more conservative, less uncertain about their lives, less likely to engage in social distancing or hygienic behaviors, and more likely to believe that lockdown restrictions were too strict ($p < .0001$ for all correlations). Overall, we highlight the need to develop interventions that decrease the discrepancy between perceived and actual risk, and implicate SVO type as a novel factor that may predict how individuals weigh competing interests and assess risk in the context of the COVID-19 pandemic.

R-35: Dynamic effects of environmental volatility on learning and choice strategy

Ethan Trepka[1], Huriye Atilgan[2], Cayla Murphy[2], Alex Kwan[2], Alireza Soltani[1]
[1]Dartmouth College, [2]Yale University School of Medicine

Characteristics of the reward environment such as reward probabilities associated with different actions and changes in these probabilities (volatility) have been shown to influence learning and decision making. The effects of volatility, however, have mainly been studied by manipulating the frequency of changes in reward probabilities (e.g. by changing the length of blocks with fixed reward probabilities). Thus, it is currently unknown whether or how higher-order statistics of the reward environment, such as predictability of block lengths, affect learning and choice behavior. To address this question, we trained mice on a two-armed bandit task with a dynamic reward schedule. Reward probabilities associated with the two actions (licking left or right) were fixed in a block of trials. Critically, block lengths were variable and depended on performance; once the animal chose the better side ten times and completed a random number of "add-on" trials, reward probabilities were reversed. We used the length of the preceding block to define the first-order statistics (volatility) of the reward environment and relative changes in block length over recent blocks to define the second-order statistics (unpredictability). While performing the task successfully, mice developed expectation of switches in reward probabilities prior to actual reversals as reflected in decreased selection of the better option over time. This effect was stronger in volatile and unpredictable environments with many add-on trials. To understand how predictive behavior influenced learning, we examined repetition index (RI), the reward-independent tendency to repeat an

action beyond chance, and the conditional entropy of reward-dependent strategy (ERDS), a measure of consistency in the use of Win-Stay and Lose-Switch strategies. The observed predictive behavior was accompanied by a slow increase in RI for the worse option and ERDS for unrewarded trials; both these effects were stronger in volatile environments. These changes in reward-independent and reward-dependent strategies allowed animals to switch to the better option faster after reversals. Moreover, after reversals mice responded more consistently to reward whereas before reversals, their response to no reward increased until it matched their response to reward. Finally, unpredictability decreased consistency in response to no reward but did not affect consistency in response to reward, indicating that unpredictability mainly impairs learning from negative feedback. Overall, our results provide evidence for how unpredictability and volatility of the environment interact to determine learning and choice strategies in a dynamic world

R-36: Neurocorrelates of gambling task risk propensity and their relationship to decision making during driving

Eric De Celles[1], Jared Boasen[1], Alexander Karran[1], Pierre-Majorique Léger[1], Sylvain Sénécal[1]
[1]HEC Montreal

Objective: Our propensity for risk potentially affects every decision of our lives, not least of which during driving. Clarifying neurocorrelates of risk-taking propensity (RTP) could permit deeper understanding of brain activity related to decision making during driving. However, despite numerous neuroimaging studies regarding driving and fatigue, investigations specifically into the relationship between risk propensity and driving decision making remain limited. We address this lacuna by exploring brain activity in cortical areas associated with risk-taking propensity during a realistic driving simulation. Methods: We recorded 64-channel electroencephalograms (EEG) from 18 healthy subjects (aged 20-39, μ =24, 6 female) during the Iowa Gambling Task (IGT) to assess RTP, and a simulated driving task which consisted in completing laps alone around a virtual race track. Risk-taking behaviour was encouraged during driving by modifying objectives and offering financial rewards. Cortical theta, alpha and beta activity during the IGT were separately regressed against IGT performance outcomes, and brain areas where activity was significantly related to performance were identified. Cortical activities in these areas during critical decision points in the driving task were then regressed against IGT performance and driving behavior telemetry. Results: IGT-based RTP was significantly positively associated with theta activity in several frontal and temporal regions of both hemispheres, notably the orbitofrontal cortices, the inferior frontal gyri, and both the superior and inferior temporal gyri. The same positive association was also found for alpha activity in left prefrontal and temporal areas. However, the significant brain activity exhibited in these regions while performing the IGT was not observed during the driving task. Furthermore, we found no significant relationship between IGT performance and driving behaviour. Conclusions: Midfrontal theta-band power has been positively associated with cognitive effort and

evaluative thinking during decision making, implying ongoing evaluation and indecisiveness in subjects with high RTP in the present study. However, IGT-based RTP had no relationship with midfrontal theta-band power during the driving task. We speculate that this discrepancy stems from the complex processing demands during driving, which may have diluted the observability of the RTP relationship. This research adds to the existing literature concerning the neural underpinnings of RTP in the IGT, while also supporting the poor generalizability of results obtained from synthetic tasks to ecologically valid contexts such as simulated driving.

R-37: Accelerated delay discounting in PTSD is associated with altered basolateral amygdala-ventromedial prefrontal cortex resting state functional connectivity

Elizabeth Olson[1], Elizabeth Wang[2], Muhammad Uddin[2], Isabelle Rosso[1]
[1]McLean Hospital/Harvard Medical School, [2]McLean Hospital

Introduction: Delay discounting (DD) is the tendency to devalue delayed rewards compared to immediate rewards. The ventromedial prefrontal cortex (VMPFC), basolateral amygdala (BLA), and nucleus accumbens (NAcc) regulate reward-based decision-making processes and play critical roles in emotional processing and self-control. Abnormal activity and resting state functional connectivity (rsFC) of these structures have been demonstrated in trauma-exposed individuals, but to date there has been limited investigation of relationships between this circuit and reward-related decision-making in trauma-exposed populations. We aimed to identify relationships between DD and rsFC between the VMPFC, BLA, and NAcc in individuals with posttraumatic stress disorder (PTSD). Methods: Participants (n = 23; 5 male, 18 female) completed clinical interviews (SCID-5, CAPS-5), questionnaires (including the Snaith-Hamilton Pleasure Scale: SHAPS), and a computerized DD fMRI paradigm. Behavioral data from this fMRI task are discussed here. Eyes-open resting state scans also were acquired. Data were preprocessed in fMRIprep and analyzed in CONN 19.c. We extracted Fisher's r-to-z values for analysis of connectivity among anatomically defined regions of interest (BLA, VMPFC, NAcc), controlling for age and gender. Results: Within individuals with PTSD, accelerated DD (higher log k) was associated with lower positive rsFC between the BLA and the VMPFC, partial r (17) = -.494, p = 0.032. DD rates were not significantly associated with PTSD total symptom severity on the CAPS-5 (or with any symptom cluster), but accelerated DD was associated with greater self-reported anhedonia on the SHAPS, partial r (17) = -.529, p = 0.020. Low positive BLA-VMPFC rsFC also was associated with higher intrusion symptoms on the CAPS-5, partial r (17) = -.550, p = 0.015, but intrusions were not significantly related to accelerated DD or anhedonia. Conclusion: Accelerated delay discounting in individuals with PTSD was associated with lower positive resting state functional connectivity between the basolateral amygdala (BLA) and ventromedial prefrontal cortex (VMPFC), and with greater self-reported anhedonia. While the imaging results do not survive stringent correction for multiple comparisons, conceptually these findings are consistent with a growing body of literature suggesting that, in individuals with psychopathology, anhedonia is associated with accelerated devaluation of delayed rewards.

Additionally, while altered VMPFC-BLA connectivity is a well-known feature of PTSD involved in aberrant fear learning and extinction, abnormalities within this circuit also may contribute to reward devaluation.

R-38: Arousal modulates dynamic decision parameters of attention in risky choice

Abdelaziz Alsharawy[1], Xiaomeng Zhang[1], Sheryl Ball[1], Alec Smith[1]
[1]Virginia Tech

Objective: We investigated the physiological and information processing mechanisms underlying the effect of incentives on risk aversion. We acquire physiological and gaze fixation data in conjunction with large changes in monetary incentives to study the neurocomputational mechanisms underlying the effect of increased incentives on risk aversion. We develop and estimate an arousal-modulated Attentional Drift Diffusion model (aADDM) of the interaction of arousal and attention during evidence accumulation. Methods: We recorded reaction time, gaze fixation, pupil dilation, pulse rate, and skin conductance while participants (N=39) chose between lotteries involving low real (1x), high hypothetical (50x), and high real (50x) stakes. Each block involved 10 choices repeated twice, with the location and order of the attributes randomly assigned, for a total of 80 choices. Results: Increased risk aversion under high real stakes was associated with changes in tonic arousal levels and in attention (as measured by gaze fixation duration). We find that the high and low attributes for available options were attended differently during both non-decision time (NDT: passive exposure phase) and decision-time (DT). Also, increased dwell time advantage on high attributes, and not low ones, was strongly associated with increased risk aversion. We estimate an attribute-based aADDM using our behavioral, eye-tracking, and physiological data. NDT and DT gaze amplified the value difference for high attributes only (multiplicative gaze bias). Trial-to-trial variation in arousal modulated this bias for NDT gaze. In other words, the high attributes' value difference is discounted steeply when attending to lower attributes, with heightened arousal further amplifying this process. No such effect was found for low attributes. Additionally, trial-to-trial increases in arousal strengthened DT additive (simple) gaze bias for low attributes and weakened it for high attributes. We also find that increased trial-to-trial arousal levels widen the decision threshold signifying higher response caution. Conclusion: We find that physiological measures and attention differ significantly when the payoffs are large and real compared to small or hypothetical stakes. We estimate a sequential sampling model that incorporates both arousal and attention. And, we find that the effect of arousal-gaze interaction differs between gaze for non-decision time and that of decision-time, with the former influencing multiplicative gaze bias and the latter influencing additive one. Our results demonstrate the value of integrating physiological and attention measures for understanding the decision-making process.

R-39: The Impact of Bipolar and Borderline Personality Disorder Traits on GPA and Impulsivity Among College Students

Jiaxuan Teng[1], Ailun Li[1], Zuzanna Tajchman[1], Iris Vilares[1]
[1]University of Minnesota

Background: Impulsivity as a multidimensional construct is commonly linked with a wide range of mental health disorders. DSM-5 states that impulsivity is a prominent factor associated with both Bipolar Disorder (BP) and Borderline Personality Disorder (BorPD). Previous research suggests that individuals with BP and BorPD are more likely to be associated with lower academic achievements. However, the relationship between these two traits and academic performance among the college population remains unclear. Study's Objective: Our study aims to measure whether non-diagnosed students possessing BP or BorPD traits will associate with poor academic performance (GPA) and high impulsivity. Additionally, the other goal of the study is to evaluate the correlation between self-report questionnaires and behavioral tasks of impulsivity. Method: Our pre-registered study included a sample of 125 college students who were measured by both behavioral tasks and self-report questionnaires. Bipolar and Borderline personality traits were measured by the Mood Disorder Questionnaire, Bipolar Spectrum Diagnostic Scale and Maclean Screening Instrument for Borderline Personality Disorder. Two behavioral tasks (Two-choice impulsivity paradigm; Go/noGo) and a self-report questionnaire (Barratt Impulsivity Scale, BIS-11) were used to measure impulsivity. Students' GPAs were also collected. Results: Both Bipolar and Borderline personality traits were positively correlated with the self-report impulsivity questionnaire (BIS-11) but not with the behavioral tasks. No significant correlation has been found between these two traits and student's GPA. Overall, the self-report impulsivity questionnaire (BIS-11) had no significant correlation with either of the behavioral tasks suggested to measure impulsivity, except the BIS motor subscale, which was slightly positively correlated with the Go/ No go task. There was no significant correlation between the two behavioral tasks. Conclusions: Our results suggest that students with bipolar and/or borderline personality traits may tend to have higher self-report impulsivity without a noticeable impact on their GPA. This result also supports the growing consensus that impulsivity describes a diverse set of processes and distinct traits.

R-40: Income shock but not negative affect increases delay discounting

Bastien Blain[1], Laura Globig[1], Tali Sharot[1]
[1]UCL

Delay discounting is sensitive to changes in the state of the individual and its environment. For example, a sudden decrease in income (i.e., income shock) increases delay discounting (Haushofer, Schunk, et al., 2013) as do negative changes to affective states (e.g., Lerner et al., 2012, but see Haushofer, Cornelisse, et al., 2013). Thus, one possibility is that the effect of income shock on delay discounting is mediated by negative affect. That is, income shocks may

lead to negative affect which alters temporal preferences. To test this, we measured income shock, affective states, and delay discounting rates in the United States (N = 200) at two points during the COVID-19 crisis (March 26-29 2020 and June 6-7 2020). We found that income shock was related to an increase in delay discounting (time 1: P = 0.003; time 2: P < 0.0001). Moreover, changes in income shock over time were related to changes in delay discounting (P = 0.02). By contrast, neither negative affective states (time 1: P = 0.21; time 2: P = 0.88), nor their change over time (P = 0.23) were associated with delay discounting. Neither did affect mediate the relationship between income shock and delay discounting. These results provide evidence for the tight link between income changes and delay discounting. However, they suggest that the increase in delay discounting is likely to result directly from liquidity constraints and not from negative alterations to affect.

R-41: Neural dissimilarity in subjective value coding of risk vs. ambiguity in vmPFC underlies more extreme ambiguity preferences

Maëlle CM Gueguen[1], Darla Bonagura[1], Melanie Ruiz[1], Gabriela Silguero[1], Anna Konova[1]
[1]Rutgers University

People are more ambiguity averse than risk averse, choosing options with known outcome probabilities over those with unknown probabilities. Clinically, ambiguity preference, but not risk preference, is found to be a marker of vulnerability to opioid use. Previously, studies have found largely overlapping neural substrates of risk and ambiguity, but this cannot explain such vast behavioral differences. Here, we hypothesize that a difference in patterns of neural activity within the value circuit could explain how preferences for ambiguity deviate from those for risk, in a diverse sample of subjects with and without opioid use disorders (OUD). 23 OUD patients (mean [SE] age, 45.6 [2.1] years; 6 women) and 21 matched control subjects (46.8 [2.8] years; 8 women) were part of a longitudinal study in a community-based treatment setting. Clinical state (anxiety, craving) and opioid reuse were assessed regularly via interview and random urine toxicology tests for up to 3 months. At a single fMRI scan, we recorded subjects' brain activity during a decision-making task manipulating separately risk and ambiguity. We used a modified utility model to parametrize the influences of risk (alpha) and ambiguity (beta) preferences on subjective value (SV). In model-based fMRI analyses, we regressed the influence of risk and ambiguity specific SV, focusing on the value circuit where SV is encoded, including the ventromedial prefrontal cortex (vmPFC). We used representational similarity analysis (RSA) to test whether more dissimilar encoding of SV translated to more divergent risk and ambiguity preferences within-subject, and differences in clinical status across subjects. Most (75.0%) individuals were ambiguity averse (median beta=0.52) and fewer (63.6%) were risk averse (median alpha=0.76), consistent with previous work. While univariate analyses revealed very similar vmPFC encoding of SV in risky and ambiguous trials, irrespective of preferences, RSA indicated more dissimilar SV patterns related to more extreme ambiguity preferences (Spearman, n=44; rho=-0.38), with ambiguity neutral subjects showing no pattern differences.

vmPFC pattern dissimilarity did not explain risk preferences across subjects (rho=-0.005) and did not relate to differences in clinical status within OUD patients. Taken together, these findings suggest risk and ambiguity preferences are instantiated via distinct representational patterns of SV encoding in vmPFC during risky decision-making. Understanding the distinct neural correlates of ambiguity and risk preferences could allow the identification of a biomarker of drug relapse vulnerability and provide more precise treatment targets.

R-42: Transcranial magnetic stimulation of the PPC leads to more rational decision making under risk

Ksenia Panidi[1], Alicia Vorobiova[1], Matteo Feurra[1], Vasily Klucharev[1]
[1]Higher School of Economics

Objective: Recent neuroeconomic research suggests that posterior parietal cortex (PPC) may play a role in risky decision-making. In the present study, we employ transcranial magnetic stimulation to explore the effects of decreased posterior parietal cortex (PPC) excitability on distinct components of risky choice such as the marginal utility of money and probability weighting. Methods: In the present study, a within-subject design is employed. Thirty five participants attended the laboratory for three sessions separated in time by 3 to 4 days (N=35). In each session, they first underwent a 40-second repetitive TMS by a perturbation continuous theta-burst (cTBS) protocol on either the right PPC or left PPC or a sham stimulation on the right PPC in a randomized and counterbalanced order. After the stimulation, participants answered a series of 85 binary lottery choice questions adapted from a widely used Holt and Laury (2002) list and presented in a randomized order purely in the gain domain. Random lottery pair design was implemented. For each experimental session we estimate hierarchical Bayesian model of rank-dependent risk preferences to obtain the distributions of model parameters on the individual and group level. We fit four different functional forms for probability weighting, including the rational case of linear probability. We found that 1-parametric Prelec function describes the data best according to the LOOIC and WAIC criteria. We then use the parameter estimates obtained for this model to analyse differences in risk preferences produced by the TMS. Results: On a group level we do not find any significant differences in either of the risk-preference parameters. However, on the individual level, we find a significant increase in the probability weighting parameter following the TMS of the left PPC relative to sham. This suggests that probability weighting function becomes more linear (less distorted) following the reduced excitability of the PPC. Additionally, we find a significant correlation of the individual effect of TMS on risk aversion with the benchmark level of risk aversion (i.e. estimated in the sham session). This means that more risk averse(seeking) participants are more likely to decrease(increase) risk aversion following TMS of the left or right PPC. Overall, these results suggest that decreasing the excitability of PPC may lead to a more rational decision making (i.e. closer to linear probability weighting and risk neutrality).

R-43: Influence of stress and anxiety on foraging decisions during Covid-19

Vani Gupta[1], Vera Ludwig[2], Dilara Berkay[2], Dhaval Bhatt[2], Michael Platt[2], Arjun Ramakrishnan[1]
[1]Indian Institute of Technology, Kanpur, [2]University of Pennsylvania

All animals including humans make foraging decisions that, in the aggregate, align with predictions of models like the Marginal Value theorem (MVT). Stress and anxiety may influence foraging decisions by changing internal states or estimates of the richness of the environment. A neural circuit involving the anterior cingulate cortex influenced by and reciprocally connected with norepinephrinergic neurons in the locus coeruleus is implicated in both foraging as well as stress and anxiety. We therefore propose that foraging decisions can serve as a readout--or biomarker--for stress and anxiety which would be useful to avoid subjective assessment through ratings, that can be confounded by lack of awareness or demand characteristics. The explosion of mental health conditions provoked by the COVID-19 pandemic makes identification of an objective biomarker even more relevant in order to identify and treat a larger number of potential patients with greater precision. In Phase 1 of the study, we recruited 161 individuals in May 2018 to play a virtual patch foraging game. We found that individuals with high trait anxiety scores stayed longer in a patch, while those with high state anxiety scores left the patch earlier than predicted by the MVT. Stress induction led to an increase in state anxiety and earlier patch-leaving decisions, but only in low trait anxious individuals. By contrast, the stress response was muted in high trait anxious individuals. A classifier trained on state anxiety, foraging behavior and EEG- and heart rate -based physiological signals could identify trait anxious individuals with ~84% accuracy. To validate and generalize this approach in a naturally stressful situation and to better understand the effects of pandemic-induced anxiety on decision making, we extended our data collection to a time point during the COVID-19 crisis (May 2020). In Phase 2, we collected data from previously-recruited individuals (May 2018 cohort, N=34) and also newly-recruited participants from two countries: 209 participants from the US and 240 individuals from India. All participants played the virtual patch foraging game and completed self-reported anxiety levels using Spielberger's State-Trait anxiety inventory. We also collected mortality rate due to COVID-19 in participant's zip code, their employment status, family situation and other measures of wellness, all of which might conceivably influence anxiety and therefore shape foraging decisions. Together, our findings suggest that behavioral variability as measured in simple foraging decisions could index stress and anxiety, which could serve as useful biomarkers in stressful situations like a pandemic.

R-44: No value-representation to rule them all: Human flexibly and adaptively tune value-learning mechanisms to learn useful representations

Tugba Altun[1], Keno Juchems[2], Andreas Jarvstad[1]
[1]City, University of London, [2]University of Oxford

Objective: Normative theory, many descriptive theories (e.g., Prospect Theory) and bare-bone reinforcement learning assume absolute-value representation. Many other theories assume representation relative to context, time and/or state. Relative values may be easier to learn, but can lead to sub-optimal choice. We propose that competing theories can be reconciled by assuming that humans flexibly adapt value-learning mechanisms to expected task demands. Methods: Participants (N=60) took on the role of manufacturing consultants for antiques (Phase 1) and vintage cars (Phase 2, counterbalanced). Each phase begun with participants learning values by sampling normally distributed market prices of pair-wise presented items. Choice without feedback followed. The Un-crossed group chose between item-pairs seen during learning. The Crossed group chose between novel pairs, made up by crossing items from different pairs. This manipulation implicitly sets expectations for Phase 2, because relative-value representation works for Un-crossed choice, but not Crossed. If mechanisms are flexibly and efficiently tuned to expectations, Un-crossed should go on to learn relative values and Crossed should learn absolute values. Value representation in Phase 2 was measured with two final surprise tasks: all-pairs 2AFC and value judgement. Results: The Crossed group learned more accurate value representations, making more high-value choices in the final 2AFC task. This difference was not due to nosier Un-crossed choices, as this group consistently preferred low-value items for diagnostic item-pairs - as predicted by relative-value encoding. Representational dissimilarity analysis was applied to participants' value judgements. Representational dissimilarity matrices (RDMs) were computed for each participant and compared to two model RDMs: relative & absolute. As predicted, the Crossed RDM was highly similar to the absolute-RDM, and the Un-crossed RDM was similar to the relative-RDM. Importantly, Crossed was better accounted for by the absolute model, and better accounted for by it than the Un-crossed group. All effect sizes were close-to-large or large. Conclusion: People do not use a single fixed value representation. Instead, they flexibly tune their value-learning mechanisms to expected task demands. When expected demands are low, people mostly learn relative values. When required, however, people readily learn absolute values.

R-45: Trade-offs between food and money or effort vary with satiety and quadratically with body mass.

Susanna Gobbi[1], Alexander Soutschek[1], Susanna Weber[1], Graf Gwendolyn[1], Daria Hinz[1], Nori Geary[1], Lori Asarian[1], Brigitte Leeners[1], Philippe Tobler[1], Todd Hare[1]
[1]University of Zurich

Objective: Rising obesity rates and associated health, economic, and social detriments are a worldwide problem. Overeating plays a key role in obesity, but the exact mechanisms behind it are still poorly understood. In this project, we analyze the predicted value of food rewards in relation to satiety and body mass index (BMI), using both money and effort as currencies. Methods: We used fMRI combined with willingness to pay and exert effort tasks to measure the value of subsequent food consumption. Specifically, participants bid either money or effort for

different snacks four times, during two menstrual cycle phases (preovulatory/ postovulatory), and levels of satiety (fasted/ fed). We investigated behavioral and neural bidding responses as a function of BMI (healthy weight: n=32, BMI: 19-25; overweight or obesity: n=34, BMI: 25-37) and satiety state. Results: We found a decrease in both willingness to pay and exert effort in the fed compared to the fasted state ($p < 0.0001$), but no linear relationship between bids and BMI or interaction between BMI and satiety. However, the bids in the two tasks depended on BMI in a quadratic way. Namely, women with relatively low or high body mass reported lower bids independently of satiety state (money: $p = 0.006$; effort: $p = 0.02$). At the neural level, the average BOLD signals were stronger in the fasted compared to the fed state in temporal and parietal regions and vice versa in the ventromedial prefrontal cortex (all $p < 0.05$, whole-brain corrected). Activity increased with higher bids in the striatum in both satiety states and currency types. In contrast, while BOLD signals in the ventromedial prefrontal cortex (VMPFC) correlated with effort bids to a similar extent in both the fasted and fed states, the VMPFC correlation with monetary bids was significantly higher in the fasted than fed state. Moreover, in both tasks we found that the relationship between bids and BOLD signals showed a U-shaped (positive quadratic) association with BMI in parietal/motor areas and inverted U-shaped (negative quadratic) association in VMPFC and temporal regions ($p < 0.05$, whole-brain corrected). Conclusions: Our data show that increased satiety decreases the willingness to consume food items independently of BMI status. However, women at the lower and higher levels of BMI within our sample reported less willingness to pay or exert effort for food rewards. These lower valuations for food could arise from multiple mechanisms, such as reduced interest in food or increased cognitive restraint towards food. Further work is required to elucidate whether the same mechanisms are at play at the lower and higher ends of the BMI range.

R-46: How much more did you like it? Understanding Utitilty Difference by Incorporating Different Decision Process Measures in Conjoint Analysis

Zhiya Zuo[1], Hongjun Ye[2], Elea Feit[2], Hasan Ayaz[3], Rajneesh Suri[2]
[1]City University of Hong Kong, [2]LeBow College of Business, Drexel University, [3]School of Biomedical Engineering, Science and Health Systems, Drexel University

Background and Objective. Choice-based conjoint analysis is a widely-used technique for assessing consumer preferences. By observing consumers' decisions in choosing between alternatives with varying attributes, consumers' preferences for the attributes can be inferred. When one alternative is chosen over the others, we infer that the decision maker perceived this option to have higher utility compared to the unchosen options. However, how large the utility difference is between the alternatives is not directly observable. Our research aims to better identify the utility difference by incorporating additional process measures including self-reported, behavioral, and neuroimaging measures. Study Design. Forty-five consumers (64.4% Female; Meanage=28.7) who were responsible to pay for their household electricity bills

participated in a conjoint study. The respondents were asked to choose between 14 pairs of hypothetical alternative pricing plans that vary on three attributes: peak rates (¢12/kWh vs. ¢15/kWh), off-peak rate (¢5/kWh vs. ¢7/kWh), and peak duration (4 hours vs 5 hours). After making each choice, participants indicated the attractiveness (AT) of each plan (1=Plan A Extremely; 7=Plan B Extremely) and the difficulty (DF) of making the choice (1=Not at all; 7=Extremely). The time participants spent on each choice task was recorded (RT). An fNIRS device (Integrated Functional Near-Infrared Spectroscopy) was employed to measure participants' prefrontal cortex activities (MO) when choosing between the pricing plans. Methods. We analyze the data with a hierarchical Bayes model, in which multinomial choices are modeled as a softmax function of a latent utility of the product attributes. To incorporate the additional process measures, we extend the model to allow AT, DF, RT, and MO to be functions of the absolute utility differences between alternatives. We use Gaussian processes to allow for flexibility in the functional form relating the process measures to the utility difference (with appropriate link functions to capture the structure of these process measures). The model is estimated using a full-information, Bayesian approach using Stan. Results. We find clear evidence that all four process measures are related to utility differences between alternatives from different perspectives. In sum, our results provide insights into how well different types of measures in conjoint studies improve estimates of consumer preferences.

R-47: Optimal Decision-Making in Metric Space

Harrison Ritz[1], Matthew Nassar[1], Michael Frank[1], Amitai Shenhav[1]
[1]Brown University

Objective The study of human decision-making has largely focused on people's choices between discrete alternatives. However, many important cost-benefit decisions are made over a metric space of outcomes, such as when planning the trajectory of a rocket to the moon. In a parallel research program, computational motor control has used normative models from control theory to explain costly, continuous bodily movement. The current project applies these control theoretic models to more abstract decision-making, in order to characterize the objective functions that people optimize during the goal-directed planning of costly choices. Methods We designed a novel task to infer how people weigh the costs and benefits of actions in a metric space. Participants simulated controlling a rover on the moon (n=30). On each trial, participants saw a 2D space with a random rover position and two random goal locations. Participants were compensated based on how close they landed their rover from both goals. Participants moved the rover using a 2D 'remote control' input space, with larger inputs incurring a larger penalty. In each block of trials, participants learned a new linear input-output mapping from the remote to the rover. Across trials, we varied the costs on large remote inputs, and the rewards earned for with proximity to each of the goals. We used inverse optimal control to find the parametrization of participants' value function that best predicted their choices. Tiling the 2D grid of possible inputs, we defined the value function at each point as the sum of rewards

(resulting goal proximity) and costs (input magnitude), using a smoothed softmax function to convert these values into choice probabilities. Results We found that an optimal control model was able to predict participants' choices with high precision, capturing the extent to which participants weighed the costs and benefits of action. Our findings also revealed interesting deviations from this reward-maximizing strategy. First, participants were not only biased away from large, costly remote inputs, but were also biased away from very small inputs, consistent with a role for input priors on choice. Second, participants' value functions were roughly quadratic, unlike the underlying reward function but consistent with popular algorithms from motor control. Conclusion Our study shows that participants use a consistent value function when making cost-benefit decisions in metric space, with promising similarities to optimal algorithms from motor control. These algorithms have the potential to provide valuable insights into human decision-making, including normative accounts of effort costs.

R-48: Assessing the dynamics of gaze and choice using webcam-based online eye-tracking

Xiaozhi Yang[1], Ian Krajbich[1]
[1]the Ohio State University

Experiments are increasingly moving online (especially during the COVID epidemic). This poses a challenge for researchers who rely on in-lab process-tracing techniques such as eye-tracking. Researchers in computer science have developed a web-based eye-tracking application (Webgazer), which has not yet received much attention in the area of judgment and decision making (Papoutsaki et al., 2016). Here, we incorporate Webgazer with the most widely used JavaScript library among behavior researchers (JsPsych), to assess the feasibility of process-tracing online. Importantly, the combination of JsPsych and Webgazer is open-source, providing straightforward access for researchers. We test our Webgazer/JsPsych combination with two decision-making studies. The first study attempts to replicate eye-tracking results from Noguchi & Stewart (2014) who studied decoy effects. However, instead of their hypothetical decisions, we adapted the stimuli and design from Trueblood et al. (2013), where subjects decide which of three rectangles has the largest area (70 trials). The second study attempts to assess gaze bias effects in binary food choice (Krajbich et al. 2010). We are using Amazon MTurk and aiming to collect 50 participants for each study. To maximize the quality of the eye-tracking data, we have implemented a number of procedures, including detailed instructions on how to optimize the webcam performance, continuous validation throughout the study, and re-calibration when necessary. In a preliminary analysis with ten subjects in the decoy task, we have achieved acceptable validation pass rates (M = 0.67, SD = 0.12). We found that most validation failures happened continuously in a small range of the trials and most were concentrated in the second half of the study. This indicates that shorter studies will likely have higher quality eye-tracking data. Behaviorally, we (so far) replicate the compromise and similarity effect, but not the attraction effect. In terms of the eye-tracking data, we find a positive relationship between relative dwell time and choice, as well as the ?gaze cascade effect? where subjects tend to be

looking at an item when they choose it. In general, we aim to provide a comprehensive analysis of choice, response time, and gaze data, as well as a ready-made template to promote online decision-making research with eye-tracking.

R-49: Characterizing approach-avoid decision-making under uncertainty in Obsessive Compulsive Disorder

Maisy Tarlow, Henry Chase[1], Simona Graur[1], Romy Frömer[2], Suzanne Haber[3], Ilya Monosov[4], Mary Phillips[1], Amitai Shenhav[2]
[1]University of Pittsburgh, [2]Brown University, [3]University of Rochester, [4]Washington University in St. Louis

Objective: Obsessive Compulsive Disorder (OCD) is characterized by persistent avoidance and a heightened sensitivity to the potential occurrence of aversive outcomes. These behaviors typically come at the cost of potential rewards. Previous work has shown that OCD patients differ from others when learning from rewarding and aversive outcomes, but has yet to examine how they weigh those outcomes when making approach-avoid decisions, particularly under uncertainty. To fill this gap, we developed a novel Probabilistic Approach Avoidance Task (PAAT) and modeled choice behavior on this task across patients and healthy controls. Method: OCD patients (N=7) and healthy controls (HC; N=10) chose between pairs of options, each involving some likelihood of a rewarding outcome (monetary gain), eliciting approach, plus some likelihood of an aversive outcome (aversive images), eliciting avoidance. Each trial had a riskier option (higher likelihoods of rewarding and aversive outcomes) and a safer option (lower of both). We varied these likelihoods across trials. After choosing, participants saw the probabilistic outcome (e.g., aversive image). Using the Hierarchical Drift Diffusion Model, we tested how subject and group level choice parameters varied with different probabilities of rewarding and aversive outcomes for the safer vs. riskier choice. Results: Overall, participants were more likely to choose the riskier option as it increased in reward likelihood and decreased in aversive outcome likelihood, and as the safer option decreased in reward likelihood and increased in aversive outcome likelihood. The influence of each likelihood on choice was captured by variability in drift rate across trials. Relative to HC, OCD patients were much more likely to choose the safer option. Modeling revealed that this finding was underpinned by two group-level differences in choice dynamics: (1) OCD patients exhibited a starting point more biased in favor of the safer option, (2) OCD patient drift rates were driven less by the reward likelihood for the riskier option. Conclusion: These preliminary findings provide novel evidence for avoidance of uncertain negative outcomes in OCD patients, and suggest that these behaviors can arise from biases in responding and in evaluating uncertain outcomes. We also validate the PAAT as a measure of approach-avoid decision-making under uncertainty, laying the groundwork for further investigation of the neural underpinnings of choosing between and subsequently awaiting the outcome of these probabilistic decisions.

R-50: Differentiating between option-wise and attribute-wise representation in multi-attribute decision making under a generative modeling framework

Qingfang Liu[1], Brandon Turner[1]
[1]The Ohio State University

It has been long debated whether people make multi-attribute choices (e.g. intertemporal choice, risky choice) by comparing different values within each attribute, or by integrating values within each option. Previous studies have suggested that eye fixation patterns provide useful insights into this debate, but they fail to explain how certain eye fixation patterns are generated and how eye fixation sequence results in eventual behavioral outcomes. To explicate the relationship between choice options, eye fixation patterns and choice outcomes, we constructed a generative model that consists of cognitive components of how people search attribute information, represent attribute values based on the searched information, and how preference is formed based on the representation. Hence, the model is able to simultaneously account for data of choice, response time and gaze fixation sequence. We found that the model simulations based on the option-wise representation, instead of the attribute-wise representation, successfully captured a variety of key summary patterns from two publicly available datasets on intertemporal choice and risky choice. We also examined whether or not the hypothesized mechanisms can be inferred from experimental data by evaluating whether current algorithms for fitting models to data can recover sensible parameter estimates. Our parameter recovery study suggested that the model parameters can be recovered using likelihood approximation techniques. We are currently working on fitting the model to experimental data. Together, the simulation result from the generative model supports the option-wise representation, not the attribute-wise representation. The generative modeling approach can be extended to many other multi-attribute tasks to model choice, response time and eye fixation together.

R-51: Neither Threat of Shock nor Acute Psychosocial Stress Affect Ambiguity Attitudes

Deshawn Sambrano[1], Arlene Lormestoire[2], Candace Raio[3], Paul Glimcher[3], Elizabeth Phelps[1]
[1]Harvard University, [2]Princeton University, [3]New York University

Economists differentiate uncertainty into two classes: risk, which has known probabilistic outcomes and ambiguity, which has unknown probabilistic outcomes. People typically find ambiguity more aversive than risk even when the ambiguous option has a higher expected value. It has been shown that the transient sympathetic arousal response to a choice predicts ambiguous but not risky decisions. Further, BOLD signal of the amygdala, a region implicated in processing emotional events, is uniquely observed to ambiguous choices. These findings suggest that emotional responses integral to the choice play a role in ambiguity attitudes.

Building on these findings, we explored the link between affect and uncertainty choices by testing whether ambiguity and risk attitudes are influenced by stress or arousal that is incidental to the choice, in two independent experiments. One experiment induced sympathetic arousal with an anticipatory threat of shock (ToS) paradigm (N = 51) and the other used a psychosocial stressor called the Trier Social Stress Test (TSST; N = 48). The manipulations' efficacy were confirmed with subjective reports of anxiety, tToS(50) = 11.08, p < .001, d = 1.48 95% CI [1.09 1.87], tTSST(46) = 5.68, p < .001, d = 1.63 95% CI [0.96 2.31] as well as with pupil dilation, tToS(7,984) = 6.89, p < .001, b = 1.12, 95% CI [0.80 1.44], and salivary cortisol, tTSST(41) = 3.75, p < .001, d = 1.23 95% CI [0.55, 1.91]. Participants made choices between a certain $5 option and either an risky or ambiguous lottery. Additionally, we used a computational model of subjective value to attain estimates of participants' preferences towards risk and ambiguity. Consistent with previous findings, participants were more averse to ambiguity than to risk. However, in contrast to our hypothesis, we found no evidence that acute stress, t(46) = .022, p = .983, d = -.006 95% CI [-.59 .58], nor arousal that is incidental to the choice, t(50) = -1.52, p = .117, d = -.103 95% CI [-.49 .28], biases ambiguity preferences. These null results were confirmed with Bayesian statistics. All Bayes Factors were less than one indicating strong evidence that these manipulations did not affect uncertainty preferences. Our study builds on the literature by testing whether affect incidental to the choice can affect decisions under uncertainty in similar ways that integral emotion does. We further find evidence that there is no casual impact of incidental arousal or stress on these types of decisions. Together these data indicate that ambiguity and risk preferences are relatively stable to minor alterations in affective context.

R-52: Brain substrates for adaptive avoidance of mental effort in second order learning

Asako Nagase[1], Keiichi Onoda[2], Toshikazu Kawagoe[3], Shuhei Yamaguchi[4], Rei Akaishi[5], Ritsuko Hanajima[1], Kenji Morita[6]
[1]Tottori University, [2]Otemon Gakuin University, [3]Rikkyo University, [4]Shimane University, [5]RIKEN, [6]The University of Tokyo

The choice mechanism based on mental effort has been recently intensively investigated. Our previous work suggested that humans adaptively learn to avoid mental effort in changing environment, and that the expected cost of mental effort for the chosen option is represented in the dorsomedial frontal cortex (dmFC)/dorsal anterior cingulate cortex (dACC) and the right anterior middle frontal gyrus (aMFG) across two different types of cognitive demand. However, the signals of expected cost in that study using a first-order learning design could be confounded with the effect of task preparation. In the present work, we used a second-order learning design, which was expected to reduce the effects of the preparation at the time of choice; participants could learn the probabilistic association between choice cues and task-type cues and the fixed association between task-type cues and either difficult or easy problems. We conducted two fMRI experiments (Exp. 1 (mental arithmetic): n = 30, Exp. 2 (spatial reasoning):

n = 28). As a result of chi-square test, most participants showed avoidance for difficult problems in both experiments (Exp. 1: 93%; Exp. 2: 92%). We fitted computational models to the choice data of avoiding participants. The best-fit model of about half of them (Exp. 1: 52%, Exp. 2: 54%) was a probabilistic model, where decisions are made based solely on the effort level of the previous trial. The best-fit model of most of the rest (Exp. 1: 37%, Exp. 2: 42%) was a reinforcement learning (RL) model, where decisions are affected by effort levels on multiple preceding trials. We conducted model-based fMRI analysis based on the RL model to data of avoiding participants. As a result, the expected cost for the chosen option was positively correlated with the activity of the dmFC/ACC, the left inferior parietal sulcus (IPS), the right IPS/medial parietal cortex (mPC), the bilateral anterior insula, and the left thalamus proper/basal ganglia in Exp. 1, and the activity of the dmFC/ACC, the right superior FG/MFG, the left IPS, and the right IPS/mPC in Exp. 2, at the time of choice-cue, regressing out the effect of decision time. The expected cost was negatively correlated with the activity of the ventromedial prefrontal cortex in Exp. 2. The dmFC/ACC, the bilateral IPS, and the right mPC were common between Exp. 1 and 2 for positive correlation, and the dmFC region included the peak coordinate of the cluster correlated with the expected cost in our previous study. These results suggest that the dmFC region represents the expected cost of mental effort for the chosen option rather than preparatory control signals during adaptive avoidance of mental effort.

R-53: Role of intelligence and learning biases in a multi-step decision-making process

Sabrine HAMROUN[1], Fabien Cerrotti[1], Mael Lebreton[1], Stefano Palminteri[1]
[1]Ecole Normale Supérieure - LNC2

Context: Cognitive science analysis of human decision making is mainly funded on two computational and neural strategies. The first, model-based strategy implies a complex computational process that maps the task structure and transitions to the earned rewards. The second, model-free strategy is a simpler strategy that maximizes a general outcome based on the learned options-rewards links. One crucial difference is that model-free strategies rely only on association between states or actions and expected rewards, while model-based strategies require more complex representations involving stimulus-stimulus associations. Experiment: The aim of our research is to assess explicit stimulus-stimulus associative structures during a multi-step reinforcement learning task. We conduct several online experiment analysis (N=106 subjects) where the two-step task was followed by post learning questions to assess the retrieval of different kinds of S-S associations. Finally, we evaluate non-verbal intelligence using a test inspired by the Raven matrices. Behavioral analysis: Our main results suggest that overall subjects chose the short-term option most of the time (65.54% of the time). Post test analysis revealed that several associations are explicitly retained, especially the option-option and the option -reward associations. State-transitions were harder to retain: 46% of correct answers for the most advantageous context compared to 69% for the less advantaged context.

Furthermore, subjects seem to remember the experienced transition, but not the whole structure.Our analysis also shows a slight correlation between intelligence and the learning of the rewards (pearson correlation=0.25, p-value=0.01) and transitions (pearson correlation=0.25, p-value=0.01). Conclusions: Our study shows that several explicit stimulus-stimulus associations are retained during multi-step reinforcement learning. However, when reward values and state-state transitions have to be concomitantly learnt, reward related associations are prioritized. Our results are relevant for the constructions of more ecological (and complex) multi-step reinforcement learning tasks.

R-54: From episodic value-based choice to temporally extended real-world behavior: Insights from a longitudinal field study of weight management goals and a minimal-narrative self model of decision making

Gina Kemp[1], Faith Li[1], Laurette Dubé[1]
[1]McGill University

The prevailing assumption in neuroeconomics is that the brain encodes representations of the expected value of stimuli and/or actions through the vmPFC, SMC and the striatum and makes a choice to maximize future expected rewards. Inter-temporal choice is often investigated using delay-discounting function which accounts for the reduction in such value over time. However, recent neuroimaging evidence sheds light on non-value signals from the hippocampus and the default network (self-related structures) that also contribute to decision making. Thus, the self becomes a key thread that connects episodic decision making with longer-term, non-value factors that are left out of classical decision models. That is why we are proposing an expanded model of decision-making that takes the self into account. We first present behavioral insights from a longitudinal field study (643 observations, 62% female) using Montreal Neighborhood Networks and Healthy Aging (MoNNETT) data panel. The study looked at the interaction between episode-level (reward sensitivity), and temporally extended self-related (prospective thinking) parameters, in relation to past and future weight-management goal success. The results show that the interplay of the two constructs in predicting future outcomes changes as a function of past goal achievements. This corroborates the above-mentioned neurological evidence, highlighting the importance of incorporating non-value, temporally extended aspects into models of decision making. In light of this, we propose a self-anchored temporally extended choice model. The self-view adopts Gallagher's notion of the self as an interaction between an episodic, executive minimal aspect and a temporally extended narrative one. The minimal self emerges from the distinction between self-versus-other with a sense of ownership and agency. The narrative dimension is temporally extended into past, present and future with a sense of continuity. Empirically, the self can be explored through studying self-related processes (SRPs), which are operations that take place during self-ideation in the resting state and are observed in the function of the default network. Our model further builds on Northoff et al. (2020) where SRPs form a Gestaltian figure-ground contrast through top-down or bottom-

up attentional processes, allowing figure SRPs to interact with the episodic context to shape decision making. In summary, choice need not be constrained to episodic functions of value and time-preference. It is rather an ongoing interaction between a multidimensional self- that is both agentic and continuous, and its environment in a dynamic process that drives adaptive behavior.

R-55: Long-term stability of idiosyncratic choice bias

Lior Lebovich[1], David Hansel[2], Yonatan Loewenstein[1]
[1]The Hebrew University of Jerusalem, [2]CNRS-UMR8002, 45 Rue des Saints Pères, 75270, Paris, France

Idiosyncratic tendency to choose one alternative over another in the absence of an identified reason is a common observation in two-alternative forced-choice experiments. Here, we show that these biases are remarkably stable over time. To this end we study the response of human participants performing two visual discrimination tasks. Repeated measures of their idiosyncratic choice bias reveal that they are strongly correlated over duration as large as 8 months. In a previous study we showed that idiosyncratic choice bias reflects microscopic heterogeneities in the connectivity of the competing neural networks that underlie choice (Lebovich et al., Nat. Hum. Behav., 2019). Imaging studies in rodents have reported conflicting evidence with respect to the stability of the cortical connectome (Mongillo et al., Curr. Opin. Neurobiol., 2017). The present results argue for the stability of the connectivity in neural networks involved in choice.

R-56: Neural and computational mechanisms for choice inconsistency

Chang-Hao Kao[1], Tiffany Dominguez[1], Caryn Lerman[2], Joseph Kable[1]
[1]University of Pennsylvania, [2]University of Southern California

Objective: People have idiosyncratic preferences in multiple choice domains: for example, they differ in their degree of temporal discounting or risk preference. However, people's choices are not always consistent with these preferences. Are these inconsistencies better characterized as mistakes or temporary fluctuations in preferences? To address these questions, we examined how brain activity, physiological arousal and reaction times were related to choice inconsistency. Methods: We re-analyzed data from a published fMRI study (n=125) where participants performed both intertemporal and risky choices. We also recruited additional participants (n=19) and collected multiple measures of physiological arousal (pupil diameter, heart rate, skin conductance) during these same tasks. In the intertemporal choices, participants chose between receiving $20 immediately and a larger amount of money after a delay. In the risky choices, participants chose between receiving $20 for certain or a larger amount of money with some probability. We first estimated participants' intertemporal and risk

preferences using a flexible non-parametric method; this makes it less likely that inconsistent choices were due to misspecification of the preference functions. Then, we identified trial-by-trial choices that were inconsistent with participant's preferences. We examined the effect of choice inconsistency on fMRI activity, physiological arousal and reaction times. Results: Inconsistent choices were correlated across time: if participants made an unlikely choice of the delayed/risky option on the previous trial, they were more likely to also do so on the current trial. Reaction times on inconsistent choices were also longer than on consistent choices on average. However, unlike errors in some cognitive tasks, inconsistent choices did not lead to slower reaction times on the following trial. In both tasks, after accounting for reaction time, choice inconsistency was associated with increased activity in dorsomedial frontal cortex (DMFC) and anterior insula (AI), and also reliably induced pupil dilation. Conclusions: These results are consistent with the idea that trial-by-trial choice inconsistencies reflect temporary fluctuations in preferences. Multiple types of random preference models predict that inconsistent choices should look like consistent choices near indifference. We find that choice inconsistency is robustly associated with increased activity in DMFC and AI, two regions that are reliably engaged during difficult choices near indifference, and this effect cannot be accounted for solely by reaction times.

S-1: Socioeconomic status is associated with grey matter volume in distributed brain regions above and beyond the effect of common genetic variants

Hyeokmoon Kweon[1], Gideon Nave[2], Martha Farah[2], Philipp Koellinger[1]
[1]VU Amsterdam, [2]University of Pennsylvania

Socioeconomic status (SES) has been known to predict cognition, behavior, and health. Recently, SES has been shown to have both genetic and neural correlates: Large-scale genome-wide association studies (GWAS) found SES to be heritable and implicated genes expressed in brain tissues. Furthermore, smaller-scale brain imaging studies found links between SES and the anatomy of numerous cortical and subcortical areas, though patterns differ across studies. Here we address two open questions about the neurobiology of SES. First, when the statistical power of MRI studies is increased by greatly enlarging sample size, which brain regions show robust SES associations? Second, given that brain structure is the product of genes and environment, which of the neuroanatomical correlates of SES can be accounted for by the known genetic correlates of SES? To this end, we conducted a large pre-registered voxel-based morphometry analysis of grey matter volumes (GMV) from 23,888 T1-weighted MRI scans of adults from the UK Biobank. We measure SES as the first two principal components (PC) of available indices of income, occupations, neighborhood, and education. These PCs were then tested for their association with voxel-level GMVs while controlling for sex, age, head size, genetic population structure (ancestry), and a number of potential confound variables. Permutation testing was used to maintain the familywise error rate at 5%. We found numerous positive associations between SES and localized GMVs throughout the brain, with the strongest signals in the frontal

pole (L,R), gyrus rectus (R), inferior occipital gyrus (R), postcentral gyrus (L), middle temporal gyrus (L), and cerebellum (L,R). Negative associations are also found, most notably in subcallosal area. To investigate the genetic contribution to these associations, we constructed a polygenic score (PGS) for SES derived from multiple GWAS results (effective N ~900K, excluding the imaged subjects) and observed the following: First, we examined the neural correlates of the SES PGS and found the strongest associations in insular, frontal, and temporal regions, some of which resemble the SES-GMV associations. Second, as implied, the magnitude of SES-GMV associations in those brain areas particularly decreased when conditioned on the PGS. Third, conditional on each other, the SES PCs and SES PGS both remain strongly associated with GMV in many regions across the brain. Taken together, our findings demonstrate robust associations between SES and GMV across the brain, which cannot be accounted for only by known genetic influences, and thus point to important environmental influence of SES on brian anatomy.

S-2: (Dis)honesty is hardwired in the brain?s functional connectome: Robust out-of-sample prediction of individual differences in cheating behavior

Sebastian Speer[1], Ale Smidts[2], Maarten Boksem[2]
[1]Erasmus University, [2]Rotterdam School of Management, Erasmus University

Cheating, manifested in financial fraud, scientific misconduct and software piracy is more prevalent than ever and inflicts immense economic costs. While it is evident that there are great individual differences in (dis)honesty, the neurocognitive determinants of this heterogeneity remain elusive. We combined neuroimaging of resting state functional connectivity (RSFC) on 99 participants with a novel experimental task (based on a paradigm proposed by Gai and Puntoni, 2017), which allows for inconspicuously measuring spontaneous and voluntary cheating, to test whether task-independent neural traits can be used to predict a person's propensity to cheat. In addition, we measured several personality traits, such as impulsivity, creativity, narcissism and greed. In our neural analysis, we used dictionary learning, which is a data-driven sparsity-based decomposition method shown to be superior to other data-driven and atlas based parcellation methods (Dadi et al., 2019), to derive 142 different functional regions. We then employed connectome-based predictive modelling (Shen et al., 2017) in combination with regularized linear regression models to identify networks that may predict cheating behaviour. There were large individual differences in (dis)honesty and we found that participants who cheated the most, also scored highest on impulsivity, which highlights the ecological validity of our paradigm. Further, we found that a network of 14 regions, including the dorsolateral prefrontal cortex and the inferior frontal gyrus, linked to cognitive control, but also the medial prefrontal cortex and temporal pole, associated with self-referential thinking, and the caudate nucleus, linked to reward processing, are predictive of a greater propensity for honesty. In a leave-one-out cross-validation analysis, we show that this neural model can reliably predict the extent to which an unseen participant will cheat. Further,

a feature importance analysis revealed that in a model combining the neural and self-reported personality measures, the neural features are significantly more important in predicting cheating than the questionnaire measures. Our findings suggest that individual differences in honesty are hard-wired in specific networks in the brain. These insights can prove instrumental in the development of more effective interventions targeted at reducing dishonesty. Moreover, the high prediction performance of our models based on stable and task-independent RSFC may be useful in the development of neuroimaging-based biomarkers for dishonesty detection.

S-3: Latent motives guide structure learning during adaptive social choice

Jeroen van Baar[1], Matthew Nassar[1], Wenning Deng[1], Oriel FeldmanHall[1]
[1]Brown University

Predicting the behavior of others is an essential part of human cognition that enables strategic social behavior (e.g., cooperation), and is impaired in multiple clinical populations. Despite its ubiquity, social prediction poses a generalization problem that remains poorly understood: We can neither assume that others will simply repeat their past behavior in new settings, nor that their future actions are entirely unrelated to the past. Here we demonstrate that humans solve this challenge using a structure learning mechanism that uncovers other people's latent, unobservable motives, such as greed and risk aversion. In three studies, participants were tasked with predicting the decisions of another player in multiple unique economic games such as the Prisoner's Dilemma. Participants achieved accurate social prediction by learning the hidden motivational structure underlying the player's actions to cooperate or defect (e.g., that greed led to defecting in some cases but cooperation in others). This motive-based abstraction enabled participants to attend to information diagnostic of the player's next move and disregard irrelevant contextual cues. Moreover, participants who successfully learned another's motives were more strategic in a subsequent competitive interaction with that player, reflecting that accurate social structure learning can lead to more optimal social behaviors. These findings demonstrate that advantageous social behavior hinges on parsimonious and generalizable mental models that leverage others' latent intentions.

S-4: Dopamine and preferences for equal payoffs

Jaime Castrellon[1], Ming Hsu[2], David Zald[3], Gregory Samanez-Larkin[1]
[1]Duke University, [2]UC Berkeley, [3]Vanderbilt University

Decisions about fairness and equity often involve tradeoffs between maximizing rewards for oneself versus others. In nonhuman animals, variation in prosocial behaviors that support cooperation and fairness mirrors differences in endogenous levels of the neurotransmitter dopamine. In particular, submissiveness and social attachment have been linked to lower levels of dopamine D2 receptors in rodents and monkeys. We hypothesized that individual differences

in dopamine function among humans may similarly reflect preferences for fairness. Here, we directly measured dopamine D2 receptor availability in a large sample of healthy adults (N=81) who underwent a positron emission tomography scan with the radiotracer [18F]fallypride. On another visit, these participants played 20 rounds of a dictator game to distribute tokens between themselves and an unknown other participant twice. Participants also distributed money as a third-party between two unknown other participants. We found evidence for associations between D2 receptor availability in the ventral striatum and amygdala and preferences for equal payoffs. Specifically, individuals with lower D2 receptor availability distributed money more fairly between themselves and others. This effect was also observed for third-party distributions between others--providing additional evidence that dopamine impacts inequity and not self-interest alone. This analysis represents the first-ever to link social decision making and in-vivo measures of dopamine function in humans. Overall, the results suggest that individual differences in dopamine function shape economic preferences and enhance our understanding of the neurobiology of social interactions.

S-5: Neural Dynamics Underlying Strategic Social Decisions

Arkady Konovalov[1], Jie Hu[1], Christian Ruff[1]
[1]University of Zurich

Strategic interactions involve complex decisions that need to integrate individual preferences, the history of play, and the decisions of interaction partners. Less complex decisions, such as food choice, risky and intertemporal choice, and static social allocations have been suggested to be governed by common choice mechanisms, since they can all be described with the drift-diffusion model (DDM). Here we sought to demonstrate that strategic decisions may also involve these choice mechanisms, and to elucidate the dynamics of these mechanisms at the neural level. We therefore combined the DDM and learning approaches to study strategic interactions in a simple repeated game, and recorded EEG to provide a dynamic neural validation for the basic concepts of this learning-DDM model: Accumulation of value for the strategic options and dynamic updates of values based on decisions of the players. Subjects (N = 40) played the standard mixed-strategy "inspection" game [1] with a real opponent (another subject), while their brain activity was measured using 128-channel EEG at 512 Hz sampling rate. On each trial, both subjects chose between letters H and T (restricted to 3 s). If the choices matched, one of the subjects received a reward, whereas in the case of mismatch, the other subject was rewarded. Each subject played the game for 400 trials with the same anonymous opponent. The behavioral data were best explained using a computational model that combined the influence learning model [2] with the standard 3-parameter DDM. In the EEG data, we identified dynamic neural correlates of various components of the fitted model, including the evidence accumulation process, the drift rate, the value of the chosen option, reward, and belief updates based on both the reward feedback the subjects' own choices. Using the clusters revealed with the model-based analysis, we found that the performance of each

subject correlated strongly with gamma-band centro-parietal coherence in the EEG signal, whereas the choice of strategy correlated with fronto-parietal gamma-band coherence. Our results offer neural validation for a model that describes choices in complex social situations as a combination of learning and DDM mechanisms, overlapping with characteristics of simple choice paradigms. Using EEG, we provide new insights into the neural processes underlying strategic choices, as well as the timing of belief updates in repeated social interactions.
Acknowledgements: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 725355 BRAINCODES).

S-7: Calculated Punishment

Fadong Chen[1], Lei Wang[1], Gideon Nave[2]
[1]Zhejiang University, [2]The Wharton School, University of Pennsylvania

Calculated Punishment [2700 characters ma] Although punishment is fundamental to the evolution of cooperation in organizations and societies, little is known about the decision processes that underlie it. Based on findings that people are faster when punishing (relative to when withholding punishment), scholars have recently proposed that punishment results from intuitive decisions. Here, we test the generality of the 'intuitive punishment' theory in an experiment of the public goods game (PGG) with punishment (N=132). Participants played 20 rounds of the PGG in groups of four, where group compositions changed every period using a stranger-matching protocol. At the punishment stage of the game, participants were shown how much each of the other group members contributed, and decided on whether to incur a cost for reducing the target's payoffs, where the cost and impact of punishment were drawn randomly and independently in each trial. We find that punishment rates and response times (RTs) in the experiment are sensitive to the cost and impact of punishment, and the contribution amounts of both the punisher and target. Moreover, people are slower to punish (relative to when withholding punishment) under some cost-benefit tradeoffs -- a finding which is at odds with the intuitive punishment hypothesis. However, the patterns in our data confirm two key predictions of sequential sampling models (SSMs): (1) an inverted-U-shaped relationship between RTs and the strength of preferences for punishing (estimated using a linear utility function that takes into account the cost of punishment, its impact, and the contributions made by the punisher and target), and (2) a positive relationship between punishment rates and the relative speed of punishment across individuals. We also confirm these predictions in the data of a previous PGG study published by Mischkowski et al. (OBHDP, 2018), which originally found support for the intuitive punishment hypothesis. To account for the above patterns, we propose a computational drift-diffusion model (DDM) that characterizes the dynamics of the punishment decision process. We demonstrate that the model improves out-of-sample predictions of PGG punishment behavior and RTs across experimental sessions, compared to computationally naïve models that rely on the same data. Our results indicate that

punishment arises from a value-based computational process that shares a common mechanism with decisions across other domains. These findings open the door for studying punishment using additional process measures, and provide a channel for influencing punishment behavior via simple interventions that target the computational

S-8: Changing Corrupt Behaviors of Power Holders via Transcranial Direct Current Stimulation (tDCS)

Yang Hu[1], Rémi Phillipe[2], Valentin Guigon[2], Sasa Zhao[2], Edmund Derrington[3], Brice Corgnet[4], James Bonaiutos[2], Jean-Claude Dreher[2]
[1]Peking University, [2]CNRS, [3]Université Claude Bernard Lyon 1, [4]EmLyon

Introduction As one of the most frequent forms of corruption, bribery takes place when a briber suborns a power holder to achieve an advantageous outcome at a cost of moral transgression. Bribery pervasively exists in governments, enterprises, and other organizations all over the world, and brings severe socio-economic consequences such as aggravating income inequality and poverty. Despite that bribery has been widely investigated in the social sciences using survey-based methods and behavioral experiments, little is known regarding the neurobiological mechanisms underpinning corrupt decision making in power holders. Recent literature in decision neuroscience has revealed a critical role of the right dorsolateral prefrontal cortex (rDLPFC) in evaluating the trade-off between personal interests and other's welfare or moral principles. Hence, the present study examined whether this brain region is causally necessary for determining whether a power holder would accept a bribe or not using tDCS. Methods One-hundred twenty healthy participants (54 females; mean age: 22.4 ± 4.4) were evenly and randomly assigned to one of the three tDCS treatment groups (i.e., anodal, cathodal, and sham). Specifically, participants in the role of a power holder decided whether to accept or reject a monetary offer (ranging from 10% to 90%) from different proposers who cheated (i.e., the bribe condition) or not (i.e., the control condition) for higher self-interest. Results We found that disrupting rDLPFC via tDCS (vs. sham) made participants more willing to accept offers in the bribe condition (vs. control), especially when the offer proportion ramped up (i.e., over 50%). Importantly, this tDCS-induced effect on corrupted behavior could not be attributed to differences in moral beliefs about bribery, sense of power over the proposer or sense of being bribed, and was not related to participants' mood and cognitive reflection ability. Furthermore, combining hierarchical Bayesian estimation and the Fehr-Schmidt inequity aversion model, computational modelling analyses revealed unveiled a causal link between the disruption of the rDLPFC and a decrease in advantageous inequity aversion (i.e., the guilt parameter) specific to bribes, which was contingent on the dispositional empathy of power holders. Conclusion The present study provides empirical evidence that perturbing rDLPFC via tDCS selectively influences corrupt decisions of a power holder, and suggests guilt as a potential cognitive mechanism supporting this causal relationship. These findings open a new window to investigate corrupt behaviors using a multi-disciplinary research approach.

S-9: The causal role of right dlPFC for norm-guided social behavior: a meta-analysis of TMS studies

Patricia Christian[1], Alexander Soutschek[1]
[1]Ludwig-Maxmilians-University Munich

Theoretical accounts ascribe right dorsolateral prefrontal cortex (rDLPFC) a crucial role for social decision making, but previous studies assessing the rDLPFC's function with transcranial magnetic stimulation (TMS) provided inconsistent evidence. While some studies suggest that rDLPFC implements social norm compliance by overriding prepotent selfish impulses, others report rDLPFC to promote (instead of inhibit) selfish choices. To decide between these conflicting accounts, we conducted a meta-analysis on studies that investigated the impact of rDLPFC TMS on social decision making. While we observed no significant effect of rDLPFC TMS across all studies, moderator analyses reveal that the rDLPFC's role in social decisions crucially depends on social context, including fairness type (proactive versus reactive) and role of decision maker (proposer versus receiver). Taken together, our results reconcile conflicting findings on the rDLPFC's role in social decision making and suggest that contextual factors determine whether the rDLPFC promotes selfish or norm-guided behavior in social interactions.

S-10: The blind leading the blind?why do we imitate when no one knows anything?

Lotem Elber-Dorozko[1], Yonatan Loewenstein[1]
[1]Hebrew University of Jerusalem

Many people knock on wood in order to ward off bad luck. But many of these people will agree that they received no direct feedback indicating that this action is indeed associated with luck. Why knock on wood then? This behavior is acquired through the imitation of others. Imitation can be useful both because others may know something that the imitator does not know ('informational influences') and because it avoids the social costs of deviating from social norms ('normative influence'). While with informational influence the imitator is swayed to believe that knocking on wood wards off bad luck, with normative influence, no such belief is necessary. How can we tell which influence leads people to knock on wood? Here we use a novel paradigm to dissect informational and normative influences on imitation, by manipulating informational influence. In our paradigm, two players play the same repeated two-alternative choice game, in which choices are associated with a monetary reward. In one condition, informational influence biases the second player to choose similarly to the first (i.e., imitate), while in other conditions it biases her to choose the opposite action, or has no effect on imitation. Normative influence is similar in all conditions. We analyze data from 121 subjects in the role of the imitators who played the game on the Amazon Mechanical Turk platform. The experiment included two block types - one with feedback about the outcome of both the

chosen and unchosen options, and one without any feedback at all. While in the first block-type participants can learn from their experience, in the second block-type such learning is not possible. We found that the magnitudes of social and informational influences are comparable. Feedback about the outcome of the actions reduces the magnitude of the total influence. Surprisingly, we found that social influence biases participants away from imitation, possibly out of a desire to be unique. Studying between-participants variability, we found that social influences are much more variable than informational ones. Taken together, these results are a first step towards a framework that allows us to predict people's behavior in scenarios where most of the given information is the choices of others.

S-11: Cognitive learning processes account for asymmetries in adaptations to new social norms

Uri Hertz[1]
[1]University of Haifa

Changes to our social settings caused by migration, cultural change or pandemics mean that we must learn and adapt to new social norms. Building on the notion that social norms provide a group of individuals with behavioural prescriptions and therefore can be inferred from individuals' behaviour, I examined how two features of the behavioural patterns of social norms, saliency and valence, affect learning and adaptation. To this end I adapted a sequential social dilemma paradigm to form a multiplayer star-harvest game. In this game four players collected stars and were permitted to sacrifice a move to zap other players. Zap outcomes were either negative or positive to others, and players could either zap each other or avoid zapping. Using this star-harvest game, I found that participants initially complied with a variety of social norms exhibited by the other players in the game. Yet after gaining experience with competitive norms, participants did not adapt their behaviour when playing with polite players. This lack of adaptation was not observed for participants moving from cooperative to competitive environments - they started behaving competitively. This asymmetry in adaptation was explained by the combined contribution of the rate of learning about others and the generalization of observations to group level. A computational model fitting procedure suggested that saliency affected learning rates as players learned more from active behaviour than from passive behaviour, while negative outcomes were more readily generalized from one player to others. These results provide a novel cognitive foundation for social norm learning and adaptation and can inform future investigations of cross-cultural differences and social adaptation.

S-12: Contextual and individual differences in altruism arise from distinct aspects of information processing

Jie Hu[1], Arkady Konovalov[1], Christian Ruff[1]
[1]University of Zurich

Cooperation and productivity in human societies depend on altruism, our ability to share resources even with genetically unrelated others. Econometric modelling suggests that altruism is not uniform and differs strongly across contexts and individuals: People on average share more when they better off than others (advantageous inequality) versus when they have less (disadvantageous inequality), but there is great heterogeneity in the individual willingness to share in these two contexts. The reason for these contextual and individual differences in altruism is unclear. Here, we uncover possible neural origins of contextual and individual variability in altruism, by combining a modified dictator game with sequential sampling models (SSMs) and EEG. Participants (N=38) made allocation decisions during both advantageous (ADV) and disadvantageous (DIS) inequality. We fitted SSMs separately for each context and compared distinct aspects of information processing (e.g., value representation, evidence accumulation, choice criteria, and biases) and the associated neural dynamics between the two contexts. Our model-based EEG results reveal three main findings. First, the same parietal signal accumulates evidence during both contexts, but the values entering this process are constructed and acted upon differently: Model estimations showed that people place greater weight on others' payoffs during ADV versus DIS, coupled with a lower decision threshold and a starting-point bias. Thus, inequality contexts affect altruistic behaviors by changing both value construction and decision criteria. Second, ERP results show that inequality is processed earlier in centroparietal sensors during ADV versus DIS, suggesting that fairness concerns are attended to differently in the two contexts. Third, individual variability in altruism related to how strongly the others' payoff was represented by centrofrontal signals in the DIS context, and to the strength of gamma-band synchronization between these centrofrontal sensors and the parietal sensors showing evidence-accumulation signals. This suggests that individual differences in altruistic behavior may originate from stronger information sharing between value-construction and evidence-accumulation regions. Our findings substantiate theoretical distinctions between different types of human altruism, by showing that contextual and individual differences formalized in econometric models relate to distinct aspects of neural information processing. Moreover, our results suggest various angles by which human altruism may be changed by behavioral and neural interventions.

S-13: The influence of confidence exchange on the quality of group-based learning

Sensen Song[1], Xiaoxue Gao[2], Zhewen He[3], Xiaolin Zhou[2]
[1]Tongji University, [2]Peking University, [3]University College London

Previous studies on group decision-making have mainly focused on either one-shot decisions or decisions that were independent of each other. Yet, a group of individuals may also make a series of decisions towards a collective goal that requires accumulative learning from previous

experience and outcomes. The mechanisms underlying such group-based learning remain unclear. Here we used a multi-round group-based learning task, in which the participant, partnered with another participant (116 participants, 58 dyads) through verbal communication, determined and learned which of the three candidate boxes contained a certain object (24 objects in total) by trial-by-trial feedbacks, and rated how confident they were about their decisions. Additionally, as a baseline control, each participant completed the same learning task by him/herself (i.e., individual learning) either before or after the group-based learning task. Results demonstrated that compared with individual learning, group-based learning exhibited both increased decision accuracy and learning rate, obtained by reinforcement learning modeling. The higher the similarity in response accuracy between the two group members is, the higher the enhancement of group-based learning compared with individual learning is. Moreover, logistic trial-by-trial regression analysis showed that in group learning conditions, choice with higher confidence made by one of group members was more likely to be selected the collective decision for the group than choice with lower confidence. Furthermore, if we always designate the choice with higher confidence or the choice with lower confidence in each trial as the collective decision for the group, the former would show significantly higher accuracy in decision than the later. These findings suggest that the similarity of learning ability between group members and the level of confidence are two crucial factors in modulating the quality of group-based learning.

S-14: Structural knowledge adaptively shapes the representation of social networks

Jae-Young Son[1], Apoorva Bhandari[1], Oriel FeldmanHall[1]
[1]Brown University

Social decision-making requires people to strategically account for the rippling effects of their actions upon others in their network. Those with a more sophisticated understanding of a network's configuration are better able to anticipate how ideas, norms, and behaviors transmit through the network. However, the number of possible configurations for how people are connected grows combinatorically with network size, and it is difficult to gain firsthand experience of all relations even in small networks. We hypothesized that people overcome these challenges using a representational strategy known as structure learning, which leverages knowledge about structure to efficiently compress information about how relationships are organized in the network, and to infer relationships that have never been observed. To test this hypothesis, we measured how subjects represented novel networks when friendships were structurally yoked to social characteristics (i.e., affiliation with an extracurricular club or sharing common interests). In Study 1 (N = 50), subjects learned about network members' friendships and social characteristics in an associative learning task, judged how socially close or distant network members were to each other, and reported their memory for friendships in the network. Using Representational Similarity Analysis (RSA), we tested whether subjects' behavioral representations of the network incorporated structural information (i.e., social

characteristics). Results reveal that representations were readily biased by structural information when it was predictive of friendships, but not when subjects knew that it was weakly predictive. In Study 2 (N = 84), we introduced a generalization task requiring subjects to infer the existence of friendships never before encountered. The accuracy of subjects' inferences was determined by their ability to use the structural information of the network members' social characteristics. Our findings demonstrate that structure learning plays a prominent role in shaping how people learn about the social world. By detecting the networks' structural properties, people can figure out how network members are interconnected to each other. Although this compression of social information introduces small inaccuracies in how people represent networks, it also enables people to draw upon structured representations to generalize beyond direct experience. Having robust structural knowledge enables people to make accurate inferences about relationships they have never observed. This capacity for flexible generalization likely aids in strategic social decision-making in the vast and constantly-evolving social world.

S-15: Predispositions and Attribute Latencies both Contribute to Preferences in Social Decision Making

Fadong Chen[1], Zhi Zhu[1], Qiang Shen[2], Ian Krajbich[3], Todd Hare[4]
[1]Zhejiang University, [2]Zhejiang University of Technology, [3]The Ohio State University, [4]University of Zurich

Social decisions involving tradeoffs between selfishness and pro-sociality are ubiquitous and important in our everyday life. Theories and computational models of social decision making usually focus on how strongly different attributes are weighted in decision making. Here, we investigate how a priori bias and the temporal dynamics of choice processes influence social decision making and explain individual differences in the effects of time pressure and delay. We use mini-dictator games in which participants (N = 117) make 300 binary decisions about how to allocate money between themselves and another participant under time-free, time-pressure (2 s limit), and time-delay (10 s wait) conditions. Using both (computer) mouse tracking and a time-varying drift diffusion model (tDDM), we find that selfish participants process their own payoffs earlier than other's payoffs, while the opposite is true for prosocial participants. The mouse-trajectory-derived relative onset times for one's own versus other's payoffs in the time-free condition correlate with participants' pro-sociality in the time pressure ($r(117)=0.722$, $p<2.2e-16$) and time delay conditions ($r(117)=0.718$, $p<2.2e-16$). Relative onset times estimated with a tDDM are highly correlated with the mouse-trajectory-derived relative onset times across all participants ($r(117)=0.679$, $p<2.2e-16$). Replicating Chen and Krajbich (2018), we find that participants' preferences changed across time conditions. That is, time pressure amplified their preferences (i.e. made participants more selfish or more prosocial), while time delay attenuated their preferences, making them less extreme. These effects of time pressure and delay were best explained by the starting point bias parameter of the tDDM (i.e. a priori,

payoff-independent bias toward selfish or prosocial outcomes), and were unrelated to the relative onset times or the relative contribution of self and other payoff magnitudes to the drift rate. Thus, time pressure and delay manipulations seem to specifically exacerbate pre-existing biases, rather than significantly altering dynamic evidence accumulation processes toward selfish or prosocial outcomes. These findings contrast with a recent report (Teoh et al. 2020), and improve our understanding and predictions of individual differences in pro-sociality.

S-16: How people attribute intention to others during strategic social decision making

Rémi PHILIPPE[1], Koosha KHALVATI[1], Rajesh P.N. RAO[2], Jean-Claude DREHER[1]
[1]CNRS, [2]University of Washington

Humans generally interact with unknown people without knowing their intentions in terms of competition and cooperation, which makes certain situations difficult to apprehend, especially when intentions of the others fluctuate over time. However, little is known about the neurocognitive mechanisms underlying inferring intentions of others in such cooperative/competitive situations. Here, we used computational modeling with fMRI in a new interactive dyadic game to investigate whether the brain dynamically tracks the implicit intention behind the actions of others and to study the neural basis of such inferences. This presumed intent is an underlying factor in strategic decision-making. Participants (n=31, 20-40 years old, 17 women) 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 between cooperative and competitive behavior (matching pennies vs hide and seek). The results showed that participants were more successful in cooperation than in competition without any information on the game mode, but did not exhibit differences in terms of switching strategies. We developed a number of computational models to capture how participants made inferences about competitive or cooperative intentions of others. Among them, the dual-intent model uses an arbitration computation defined as the difference in reliability of the cooperative and competitive modes on a trial-by-try basis to weight decisions and adjust behavior. Using a Variational Bayesian comparison between models, we found that the dual-intent model outperforms the others, revealing a population bias towards the competitive mode. Unlike the actual interaction mode, the arbitration computation predicts the expected intent of the other and the change in strategy. That is, the more reliable the competitive framework is compared to the cooperative framework, the more volatile and unpredictable participants tend to be. The fMRI analyses revealed that the difference in reliability (cooperative minus competitive) positively correlated with bilateral ventral striatum and caudate nucleus activity. Finally, a competitive bias involved areas engaged with theory of mind (e.g. left dlPFC, left Angular gurus). Taken together, these results indicate that the brain dynamically monitors the intentions underlying the actions of others in terms of competition and cooperation and adapts behavior to these intentions.

S-20: Attention to luck versus effort information underlies self-serving biases in allocation decisions

Dianna Amasino[1], Davide Pace[1], Joël van der Weele[1]
[1]University of Amsterdam

Ideas about how to fairly divide joint earnings varies widely from equal splits to dividing based on effort to incorporating random luck as well. Previous research has found evidence for self-serving biases in choosing the fairness criteria for division; in particular, workers randomly assigned a higher pay rate keep more of the joint earnings regardless of effort (Konow, 2000). Here, we investigate attention as a mechanism underlying this bias. We explore whether a random advantage in pay rate influences information search patterns to effort vs. luck information. We then examine if variation in attention impacts allocation decisions and reinforces this self-serving bias. Participants and passive recipients completed a variety of real effort tasks to produce a surplus. Participants were split between high (N=100) or low (N=100) pay rates per correct answer and paired with recipients (N=300) with opposite pay rates. On each trial (20 total), participants had 6s during which they could reveal their own and the paired recipient's effort (correct answers) and monetary contributions (pay rate x correct answers) in MouseLabWeb before deciding how to divide the total surplus. In accordance with previous findings, participants with high pay kept significantly more (~12%) of the surplus for themselves than those with low pay, controlling for effort. Information seeking also diverged between conditions with a difference in dwell time on effort versus monetary information. This difference developed over time as indicated by a significant interaction between condition and trial number in predicting attention, with increasing time spent on monetary information by high-pay participants and effort information by low-pay participants. Finally, differences in attention also impact divisions. There is a significant interaction between condition and attention in predicting divisions such that longer dwell time on effort versus monetary information decreases the amount kept by high-pay participants but increases the amount kept by low-pay participants. Therefore, not only do participants look more at advantageous information, but this difference in attention impacts subsequent allocation decisions. We show that a randomly assigned high-pay rate leads to a self-serving allocation and extend this finding to include the role of attention. Our attention data reveals that participants increasingly sought out diverging information depending on their condition. Furthermore, we find that this divergence in attention impacts allocation decisions, suggesting that this attentional bias may help justify self-serving conceptions of fairness.

S-21: Functioning of resting-state networks predicts the propensity to selfish behavior in the Iterated Prisoner?s Dilemma

Krzysztof Bielski[1], Sylwia Adamus[1], Iwona Szatkowska[1]
[1]Nencki Institute of Experimental Biology PAS

Iterated Prisoner's Dilemma (IPD) is a theoretical model of many social and political phenomena such as free - riding or climate-change policies negotiations. In this game, choosing defection usually leads to more individually beneficial outcomes, while selecting cooperation can results in collectively beneficial solution. For that reason, it is important to recognize determinants and prerequisites of selfish behavior leading to exploitation of natural and human resources. The general level of trust or propensities toward cheating on others might be encoded in the intrinsic brain activity as an individual characteristic of each person. Such the activity reveals the stable patterns of synchronization known as resting-state networks (RSNs) which can predict personality traits and other features observed in behavioral performance or cognitive abilities. Thus, the main aim of this research was to explore whether the functional connectivity of the RSNs might be a good predictor of prosocial or proself attitudes in behavior of people facing problem of the IPD structure. To address this issue, we recruited 30 healthy subjects who played the IPD with computer algorithm as a counter partner (being convinced that they play with another subject). Before the game, they were scanned using the Siemens Trio 3T during the resting-state for 15 minutes (eyes-open, TR=1.25 s., voxel-size = 3 x 3 x 3 mm., 586 volumes). After the game, we computed several behavioral measures observed during the game e.g. proportion of cooperative choices, percentage of reciprocated cooperation etc. Neuroimaging data were decomposed into 30 components using group-level Independent Component Analysis (MELODIC ICA) which were labeled as well-known RSNs such as Default Mode Network (DMN), Frontoparietal, Dorsal Attention etc. Then, we tested how connectivity within those networks predicted behavioral performance indicators by applying the multivariate linear regression model with the leave-one-out cross validation procedure. It showed that connectivities within Frontoparietal Right and Dorsal Attention Networks were positively correlated with unreciprocated cooperation in the IPD. The result suggests that the general propensity to selfish behavior might be embedded in the functioning of these networks. Supported by the Ministry of Science and Higher Education of Poland: 0135/DIA/2017/46.

S-22: Reinforcement by evidence of correct choice in a social decision-making task

Daniel Parr[1], Alexandre Dombrovski[2], Michael Hallquist[1]
[1]University of North Carolina at Chapel Hill, [2]University of Pittsburg

Background: Just as rewarded actions tend to be repeated, unchosen actions may be reinforced by evidence that the action would have been rewarded. Such learning from counterfactuals may be supported by a model-based decision process in which experiences beyond past choice outcomes inform predictions of potential actions' consequences. An alternative account -- reflected in theories of regret-based decision-making (e.g., Coricelli, Dolan & Sirigu, 2005), fictive error processing (Lohrenz et al., 2007; Kuhnen & Knutson, 2005), and elsewhere -- is that

choosing correctly is intrinsically rewarding, permitting evidence that unchosen options were inferior to reinforce choice. In this study, we investigate the reinforcement value of choice feedback in a social decision-making task. Methods: Sixty-four psychologically healthy participants completed a trust game during an fMRI scan. On each trial, subjects chose between keeping $1 and sharing it with a "trustee." If the trustee also shared, subjects received $1.50; otherwise they received $0. When the subject kept ("keep trials"), they were told the trustee's choice, and thus learned what the outcome of sharing would have been. We fit subjects' choices using six different reinforcement learning models: two models in which reinforcement was based on monetary outcomes and four "correctness" models in which reinforcement was based on the correctness of choices (i.e., whether subjects' choices matched trustees'). Further, we constructed multi-level models in which subjects' decisions were predicted by trustees' past choices. Results: A multi-level model that included an effect for trustees' decisions on keep trials better predicted subjects' choices than a model which did not, confirming that subjects learned from counterfactual information ($X^2[2] = 34.4$, $p < .0001$). Bayesian model comparison favored a "correctness" model, but differences between model fits were not significant (exceedance probability = .999, protected e.p. = .522). For keep trials, but not share trials, we entered the reward prediction error (RPE) signal from the favored correctness model into a voxelwise GLM analysis. Any reinforcement learning model that tracks monetary values predicts no RPE-related activation on keep trials. However, this analysis revealed robust activation in bilateral ventral striatum ($p < .001$), a region closely associated with RPE representation (Chase et al., 2015). Conclusions: In a social decision-making task, striatal learning signals tracked the correctness of choices, not simply the rewards received. Such signals may support learning from counterfactual information in social contexts.

S-24: Computational model of Minnesota Trust Game explains suspiciousness towards partner

Rebecca Kazinka[1], Iris Vilares[1], Angus MacDonald[1], A. David Redish[1]
[1]University of Minnesota

Our goal was to model spite sensitivity, i.e. worry that another person will intentionally incur a loss to ensure you do as well, by computationally analyzing an economic game designed to index persecutory ideation called the Minnesota Trust Game (MTG). To this end, we modified the Fehr-Schmidt (1999) inequity aversion model by allowing their guilt parameter to take on negative values, which provided a measure of spite; Additional risk parameters were included. 243 undergraduates completed the MTG, which is a modified version of the trust game. Participants decide if they want to take a smaller, certain amount of money ($10 for both) or trust a partner to choose between a fair outcome ($20 for both) or an unfair outcome consisting of a temptation payout for the partner and an adverse payout for the participant. In the Rational Mistrust condition, the temptation for partners was $25, while in the Suspiciousness condition it was $15. This condition measured spite sensitivity, as the player

might not trust if they believed the partner was willing to lose money just so the participant would as well. The adverse payoff for the player varied between -$5 and $15. Participants played against both a human partner and a random (50-50) partner (a coin). Parameter estimates were compared to individual differences in behavioral and personality measures. Behaviorally, most individuals trusted less as the adverse payoff decreased below $10. However, in the Suspiciousness condition, responses varied in the willingness to trust the partner, such that lower suspicious beliefs were associated with greater trust. The best model included envy of the partner's greater outcome and 3 additional parameters; separate risk aversion parameters were fitted for each partner type (random vs. human), and we also modeled the estimation of the partner's guilt of additional winnings, allowing for it to take negative values representing beliefs of a spiteful partner. The best model only estimated the human partner's guilt, while the random partner (a coin) was assumed to have a 50-50 chance of picking either option. Our modification of the Fehr-Schmidt model improved estimation of parameters of the MTG. Importantly, we isolated the estimation of partner's guilt, which allowed a direct measurement of the sensitivity to spite. This measure was highly correlated with the behavior in the Suspiciousness condition most associated with persecutory ideation. This additional condition, unique to the MTG, can be modeled to isolate spite sensitivity, suggesting that spite sensitivity is separate from inequity aversion or risk aversion and provides a means of quantifying persecutory ideation.

S-25: Financial Exploitation in Older Adults: Characterizing the Role of Sociodemographic Factors, Cognition, and Social Decision Making

Srikar Katta[1], Katherine Hackett[1], Johanna Jarcho[1], Tania Giovannetti[1], Dominic Fareri[2], David Smith[1]
[1]Temple University, [2]Adelphi University

Older adults are often the targets of financial exploitation. While most cases of financial exploitation are perpetrated by strangers, perpetration by individuals within the victim's social network (e.g., friends and family) is common. Although these observations suggest a social nature to financial exploitation, very little is known about how older adults integrate information from the social domain to inform economic decision making. To address this issue, we conducted a pre-registered survey study (https://osf.io/fcrus) that targeted adults (N = 625; ages 50-91 years; Mage = 62.24 years) in the Philadelphia, PA region. We assessed risk for financial exploitation using the Older Adult Financial Exploitation Measure (OAFEM), a 30-item questionnaire on exploitation experiences and risk factors over the past year. Our survey also included measures of social decision making (Trust Game and Ultimatum Game) and self-report measures of gullibility (with subscales of persuadability and insensitivity to trustworthiness cues), the need to belong, social support, cognitive decline, and health (physical and mental). We first examined the links between willingness to trust and risk for financial exploitation. We found that the relation between risk for financial exploitation and insensitivity to

trustworthiness cues was moderated by physical health (F(1,621) = 6.09, p = 0.013). Specifically, those with greater insensitivity to trustworthiness and poorer health were at the highest risk for exploitation. We also examined whether financial risk was associated with willingness to trust friends relative to strangers. We found older adults with the highest OAFEM scores also exhibited a higher need to belong and had a reduced capacity to differentially trust friends relative to strangers (F(1,621) = 6.01, p = 0.014). Next, we examined the links between social support and risk for financial exploitation. Although our analyses supported the prediction that individuals with less social support would be at greatest risk for financial exploitation, this relationship was moderated by physical health (F(1,621) = 4.25, p = 0.039), socioeconomic status (F(1,621) = 6.42, p = 0.0115), and cognitive decline (F(1,621) = 5.06, p = 0.025). Specifically, social support was lower among those at high risk for financial exploitation who also had poorer health, lower socioeconomic status, and higher cognitive decline. Taken together, these results support the idea that risk for financial exploitation is associated with altered social decision-making processes and poor social support, and these relations are moderated by health, socioeconomic status, and cognitive decline.

S-26: Older Adults Exhibit Enhanced Connectivity between Caudate and Default Mode Network During Shared Reward Processing

Dominic Fareri[1], Katherine Hackett[2], Tania Giovannetti[2], David Smith[2]
[1]Adelphi University, [2]Temple University

Social connection supports our overall well-being throughout the lifespan, in part by satisfying a fundamental need to belong. We have previously shown in healthy young adults that rewards experienced with close friends (vs. strangers) elicit stronger activation in the ventral striatum. Normative aging is associated with an increased importance placed on present-oriented socioemotional goals, alterations in corticostriatal function, and increased susceptibility to financial fraud. Yet, it remains unclear whether the integration of social and reward value differs between older and younger adults. We conducted a block design version of a reward processing paradigm in which young (ages 18-30; n = 24) and older (ages 65-80; n = 17) adult participants had the opportunity to share monetary rewards and losses with a same-sex close friend, a stranger or a computer (one partner per block) while undergoing fMRI. Each block of the task was also either a reward block (75% rewards) or a loss block (75% losses). On each trial, participants guessed whether the value of a card was higher or lower than the number 5; correct guesses resulted in a shared gain of $10, incorrect guesses resulted in a shared loss of $5. A repeated measures ANOVA on post-session subjective ratings of shared outcomes revealed a significant partner x age group interaction (F(2,76) = 3.97, p < 0.025): relative to younger adults, older adults rated rewards and losses shared with social partners more positively than those shared with the computer. Whole brain analyses of fMRI data (cluster level corrected, z > 3.1) revealed significant age-related differences in temporoparietal junction (TPJ) responses during shared reward relative to loss blocks: TPJ responses in younger adults

were modulated by social closeness with a partner, whereas older adults demonstrated no differentiation in the TPJ with respect to social closeness. Finally, given that the default-mode network (DMN) has been associated with reward processing and social cognition, we examined whether any brain regions exhibited age-related differences in task-based connectivity with the DMN during shared reward processing. Strikingly, we found that older adults (relative to younger adults) showed enhanced DMN-caudate connectivity as a function of partner. In sum, our findings suggest alterations in the integration of social and reward-related information in older adults, which may have implications for understanding the increased rates of financial fraud often observed in the elderly.

S-27: Intentions and actions: Punishment and compensation decisions in response to social norm violations

Mirre Stallen[1], Anabela Cantiani[2], Alan Sanfey[2]
[1]Leiden University, [2]Radboud University

Previous research has shown that people are willing to incur often quite significant costs to punish the perpetrator of a social norm violation, both when they are directly affected by the injustice themselves, but also when they have observed the violation occurring to someone else. Further, in this 'third-party' observer role, people typically prefer to punish a perpetrator as compared to compensating the victim (Stallen et al., 2018). An important open question, however, is to what extent the decisions to punish and compensate depend on the intention of the perpetrator, as opposed to the eventual outcome of their action. If someone inadvertently hurts someone else, should they be held responsible for the unintended outcomes? Similarly, if someone attempts a norm violation but fate intervenes to spare the potential victim, should the perpetrator be punished anyway? The current experiment was designed to answer these questions. Methods: Participants (n=102; within-subject design) performed an online version of the Justice Game (Stallen et al., 2018) on the MTurk platform. Participants were either treated unfairly themselves and had the option to punish the perpetrator in response (second-party punishment condition), or they observed a third-party being treated unfairly and could either punish the perpetrator (third-party punishment condition), or compensate the victim (third-party compensation condition). On each trial, participants saw both the intention of the perpetrator as well as the actual decision outcome. On half of the trials intentions and outcomes were aligned - the perpetrator's intentions, either to violate or not, were carried out as planned). On the other half of the trials, however, intentions and actions were not aligned - perpetrators either accidentally harmed their partner or failed to hurt them as intended. Results: We found that knowing the perpetrator's intention significantly affected participants' willingness to punish, but that this response depended on whether they could punish the perpetrator or compensate the victim. Specifically, in punishment decisions, knowledge about the intention of the perpetrator had a greater impact than the actual harm done, both to participants' themselves or to a third-party. However, when compensating a victim,

participants' decisions depended more on the actual harm done than on the intention of the perpetrator. Conclusion: Research on altruistic punishment has typically focused on the perpetrators' actions. This study expands on this work by demonstrating the importance of additionally understanding the intentions of the perpetrator in assessing punishment and compensation decisions.

S-28: Neurocomputational mechanisms distinguishing Rawlsian and utilitarian moral preferences

Yoonseo Zoh[1], Soyeon Kim[2], Hackjin Kim[3], Molly Crockett[1], Woo-Young Ahn[2]
[1]Yale University, [2]Seoul National University, [3]Korea University

A common ethical dilemma pits the welfare of a single person against the greater good, as in the classic "trolley problem" where individuals are sacrificed to save many others. Such dilemmas highlight a tension between utilitarian principles that require prioritizing the greater good, and Rawlsian principles that rather advocate minimizing harm to the worst-off individual. While reasoning about hypothetical versions of this dilemma has been extensively studied, little is known about how people make actual decisions that pit two conflicting moral principles against one another and the neural mechanism involved in such decisions. Here, we investigated the neurocomputational mechanisms distinguishing Rawlsian and utilitarian moral preferences. To this end, we deployed a novel fMRI paradigm where participants (N = 52) had to repeatedly decide how to allocate pain between a single individual ("the one") and a group of 3-4 individuals ("the group"). Across trials, we parametrically varied the amounts of pain delivered to the one vs. the group. Crucially, on some trials, minimizing the total amount of pain required delivering more pain to the one than to each member of the group, thus pitting the "minimax" Ralwsian strategy against the utilitarian strategy of minimizing the overall amount of pain. We used computational modeling to capture how people compute the value of others' pain and weight the pain of the one against overall pain. Using model-based fMRI, we measured the neural encoding of decision values as well as the objective amounts of pain delivered to the one versus the group, and related these neural signals to individual differences in computational model parameters. Behaviorally, participants overall favored the Rawlsian strategy, preferring to impose more pain overall so as not to unfairly target the one individual. Computational modeling revealed two dissociable dimensions of individual variability in moral choices: a parameter alpha captured the weight reflecting the minimax principle and while phi reflected the subjective threshold of acceptable amount of pain to impose on one person. fMRI revealed that moral decisions pitting Rawlsian against utilitarian principles recruited the mentalizing and valuation networks. Across participants, Rawlsian preferences were associated with stronger responses in the mentalizing network and weaker responses in vmPFC to the pain of the one relative to the pain of the group. Our results reveal the neurocomputational mechanisms guiding tradeoffs between the welfare of one versus all, and highlight the role of the mentalizing network in shaping Rawlsian preferences to protect the worst off individuals.

S-29: Patient's Response to Therapist's Reciprocity in Trust Game Predicts Psychotherapeutic Relationship

Jihan Ryu[1], Matthew Heflin[1], Soojung Na[1], Yi Luo[2], Rebeccah Bortz[1], Xiaosi Gu[1]
[1]Icahn School of Medicine at Mount Sinai, [2]Virginia Tech Carilion Research Institute

Background: Strong patient-therapist alliance is a key variable to successful outcomes in psychotherapy treatment. However, how the construct of alliance manifests in the quantitative behavioral interaction between two individuals is largely unknown. Objective: The aim of the current study is to develop a novel application of Trust Game (King-Casas, 2005) paradigm in psychotherapy dyads to identify quantitative markers descriptive of therapeutic alliance. Methods: This is an observational cross-sectional study of therapy pairs in the academic teaching hospital-based outpatient psychiatric clinic in Manhattan. Individual perception of therapeutic alliance was measured using Working Alliance Inventory - Short Form (A. O. Hovarth, 2006) after a psychotherapy session, and dyad alliance score and difference between patient and therapist were constructed as clinical variables. Patients and therapists then separately played a trustee in Trust Game after being instructed to mentalize the investor as their therapy partner. Investor's reciprocity was calculated for each round. We first regressed trustee's changes in repayment on investor's reciprocity, and then regressed therapeutic alliance on the individual coefficients, as well as average repayment fraction, a proxy index for trustworthiness of investor across the rounds. Results: In the Trust Game played by n=9 pairs (5 psychodynamic psychotherapy, diagnosis = 7 personality disorder), investor's reciprocity linearly predicted the changes in patients' repayment behavior (b=0.49, df=58, p= 9.2e-05), but not therapist's repayment behavior (b=-0.15, df=65, p = 0.28). Patients' individual sensitivity to reciprocity positively predicted the individual perception of alliance (b=0.65, df=7, p = 0.056, R-squared=0.43) and negatively predicted dyad alliance score differences (b=-0.67, df=7, p = 0.047, R-squared=0.45). Neither the patients' nor therapists' average repayment fraction across the rounds predicted the individual (Patient b=0.63, p = 0.07; Therapist b=0.09, p = 0.83; df = 7) or dyadic alliance scores (Patient b=0.33, p = 0.39; Therapist b=0.06, p = 0.89) Conclusions: Unlike therapists, patients in real-life treatment settings increased repayments in response to positive reciprocity and reduced repayments in response to negative reciprocity, indicating their different social norms in the clinical relationship. Patients' responsiveness to the reciprocity shown by therapists, not the amount of repayment towards them, explained therapeutic alliance. Objective marker to describe behavioral changes in the interpersonal context can be a novel approach to assess therapy patients' therapeutic alliance.

S-30: Modeling Dynamic, Competitive Decision-Making in Social and Non-Social Contexts

Kelsey McDonald[1], John Pearson[1], Scott Huettel[1]
[1]Duke University

Understanding how humans make real-world decisions in dynamic environments is a key goal of decision neuroscience. Neuroscientists and psychologists, however, have largely used discrete experimental paradigms aimed at constraining behavioral complexity (e.g. choices in discrete-play games, such as a bandit task) to study how humans make decisions. Here, we describe a series of projects that investigate how humans behave and make goal-directed decisions in both social and nonsocial dynamic environments. Our computational approach developed Bayesian nonparametric models (i.e. Gaussian Processes) that characterize individual strategic behavior; these models allow quantification of the instantaneous dynamic coupling between agents, mirroring the dynamic, coevolving decisions among interacting agents that humans make in real-world decisions. We subsequently applied these models to fMRI data while humans played a competitive real-time video game against both human and computer opponents. We found that the dorsolateral prefrontal cortex displayed selective activation when the subject's actions were highly sensitive to the opponent's actions, whereas activation in the dorsomedial prefrontal cortex increased proportionally to the advantageous timing of actions to defeat one's opponent. Moreover, the temporoparietal junction tracked both of these behavioral quantities as well as opponent social identity, indicating a more general role in monitoring other social agents. These neuroimaging results suggest an overlap in brain regions shown to play an important role in social cognition and value-based decision-making and regions that contribute to the goal-directed tracking of value of social actions in dynamic, multi-agent contexts.