

## P1-G-26 - Examining the associations between socioeconomic status and temporal discounting across monetary, social, and health rewards

Christian Benitez<sup>1</sup>, Michelle Mahecha Perez,<sup>1</sup> Joe Kable<sup>1</sup>

<sup>1</sup>University of Pennsylvania

**Objective:** Several studies have shown people from lower SES backgrounds discount the future to a greater extent, and multiple theoretical models have been proposed to account for this phenomenon. One account focuses on the increased financial constraints that lower SES individuals face, including limited savings and access to borrowing options. Consequently, this leads to a prioritization of short-term goals to meet daily financial needs. Another perspective emphasizes lower SES individuals' exposure to early life stressors, leading them to prefer options that minimize the cost of uncertain and unstable future environments. Discriminating between these models is challenging because prior research has solely focused on monetary rewards, while the accounts differ primarily in the extent to which they would predict higher discounting across a broad domain of potential rewards. To address this gap in the literature, this pre-registered study examines the relationship between SES and intertemporal discounting across monetary, social, and health rewards.

**Methods:** 1006 online participants completed three intertemporal choice tasks followed by a series of self-report current and childhood socioeconomic measures.

**Results:** An ANCOVA with continuous SES as a covariate on decision preferences revealed a main effect of the reward domain ( $F(2,1989) = 88.61, p < .001$ ) and current SES ( $F(1,1989) = 27.50, p < .001$ ). Replicating previous findings, we find that lower current SES ( $r = -0.162, p < 0.001$ ) was predictive of preferences for immediate monetary rewards. We also find a similar relationship for current SES and health rewards, ( $r = -0.162, p < 0.001$ ). However, we did not observe a relationship between SES and temporal discounting of social rewards. Overall, this suggests that the relationships between SES and time preference is stronger for money ( $z = -3.71, p < .001$ ) & health ( $z = -3.06, p < .005$ ) than for social rewards.

**Conclusion:** Our results demonstrate that individuals from lower socioeconomic backgrounds exhibit a preference for immediacy for health and monetary rewards that is not observed for social rewards. Importantly, this suggests that lower SES is not associated with a universal preference for immediacy, but rather reward-specific. One possible explanation is that SES places specific constraints on monetary and health choices, which shape discounting behavior. Future research can expand upon these findings by examining health discounting in societies outside of the US where health care is universally provided. This can provide further insight into the cultural and systemic factors that shape temporal discounting.

## P1-G-27 - Quantifying a task-invariant Bayesian prior for active avoidance: A pilot study

Tobias Granwald<sup>1</sup>, Peter Dayan<sup>2</sup>, Máté Lengyel<sup>3</sup>, Marc Guitart Masip<sup>1</sup>

<sup>1</sup>Karolinska Institutet, <sup>2</sup>Max Planck Institute for Biological Cybernetics & University of Tübingen, <sup>3</sup>University of Cambridge & Central European University

In a Bayesian framework, decision-making involves combining information available at the time of the decision with prior experiences or expectations. The prior is especially important since available information can often be ambiguous or incomplete. As such, priors are thought to have a profound impact on behaviour. For example, the generalised passive behaviour in patients suffering from learned helplessness and depression has been conceptualised as resulting from a pessimistic prior implying the expectation of failure when performing an action.

When measuring these priors in the lab, it is important that they generalise across different situations. Otherwise, the priors captured may only reflect task-specific assumptions that do not apply more widely. Task-invariant perceptual priors have been measured using cognitive tasks, however, task-invariant priors for affective decision-making have not been quantified.

To quantify a generalizable prior that actions will be successful, we administered two differently framed decision-making tasks. In the tasks, 54 participants made repeated decisions as to whether to take an active action or not. The passive choice was associated with a sure loss or missing out on a potential reward while the active action was associated with a cost and a probability to avoid the negative outcome. Thereby, the value of the passive choice is fully transparent. We then used computational modeling to characterize each participant's prior expectation that an active choice will result in avoiding a negative outcome. Parameter and model recovery for all models were satisfactory.

We found that a Bayesian prior model fit better than 6 alternative models (Bayesian model with a beta distribution prior vs the best-fit model without Bayesian updating: Task 1, model frequency (MF)=.648, exceedance probability (EP)=.985; Task 2, MF=.658, EP=.990). Furthermore, the means and variances of the priors correlated across tasks (M:  $r = .407, p = .003$ , var:  $r = .271, p = .047$ ). When explicitly modelled, slightly more participants were fitted by a model with a common prior across the tasks than a model with one prior for each task (MF=.535, EP=.692). Furthermore, the mean of this common prior was moderately reliable one week later ( $N=43, ICC(2,1) = .569$ ). The common prior also correlated with positive mood (M:  $r = .244, p = .075$ , var:  $r = -.310, p = .022$ ).

In conclusion, our data show that it is possible to use decision-making tasks to quantify prior beliefs that an active choice will result in avoiding a negative outcome. This method will be used to quantify these priors in clinical populations. This may provide insights into learned helplessness and depression.

## P1-H-28 - Context-dependent P2 plasticity dynamics in the monetary incentive delay task

Anastasia Grigoreva<sup>1</sup>, Alla Kondratenko<sup>1</sup>, Aleksei Gorin<sup>1</sup>

<sup>1</sup>HSE University

Recent advances in neuroeconomics and decision-making research have emphasized the neural mechanisms underlying feedback processing in reinforcement learning. Understanding how humans process feedback, predictions, and prediction errors, as well as their relationship to reinforcement magnitude, is crucial both for elucidating the mechanisms driving human economic behavior and for informing the development of new models of neuroplasticity and learning-related brain reorganization. In reinforcement learning, the value associated with predicted outcomes is also of key importance. It is postulated that reward or loss prediction errors generate adaptive signals that drive behavioral changes.

The aim of our study was to investigate learning-associated plastic changes by examining alterations in evoked potentials in response to auditory cues during one of the most used empirical trials in neuroeconomics – monetary incentive delay (MID) task. To do so, we recorded EEG during the loss version MID task performance and assessed how responses to auditory signals reflected potential

monetary loss evolved. We focused on the P200 evoked activity component, which has been linked to improvements in sound perception and performance in numerous auditory tasks.

We analyzed data from 32 participants who completed the experiment over two days. We compared event-related potentials (ERPs) between the first and second day of the task and across different contexts, anticipating significant differences in P200 amplitude. A variance analysis of the P2 component amplitude revealed the significance of Day factor ( $F(1, 31) = 12.711, p = 0.00120, \hat{1}\cdot p2 = 0.29$ ) and Day\*Context factors interaction ( $F(2, 62) = 4.7554, p = 0.01198, \hat{1}\cdot p2 = 0.13$ ). Post-hoc analysis using the Bonferroni method highlighted significant differences in contexts involving high losses and distinct losses.

Our findings demonstrate that associating a previously neutral auditory stimulus with a specific monetary loss can induce plastic changes, manifesting as context-dependent enhancement of auditory evoked potentials amplitudes. These results contribute to existing models of learning and decision-making and add to our understanding of the neural mechanisms underlying reinforcement learning.

### P1-H-29 - Accounting for Multiscale Processing in Adaptive Real-World Decision-Making via the Hippocampus

Dhruv Mehrotra<sup>1</sup>, Laurette Dube<sup>2</sup>, Goktug Bender<sup>2</sup>

<sup>1</sup>Integrated Program in Neuroscience, McGill University, Canada, <sup>2</sup>McGill University

For adaptive decision-making in diverse and changing real-world contexts, the brain needs to allow past information to influence current processing over multiple timescales, enabling complex and reciprocal dynamics between the two underlying reinforcement learning strategies: model-free (MF) - a reflection of the past - and model-based (MB) - a deliberation on causal environment-behavior structure. Yet, while the role of the hippocampus in memory and spatial learning is well-known, its interactions with the striatum with regards to the “here and now” decision-making is just starting to be explored. This work aims to better appreciate the role of the hippocampus in decision-making. Typically, humans rarely act in a pure MF or MB manner. We therefore review literature in reinforcement learning that examines the successor representation (i.e., an expectation of future states from a given starting state) as a means to bridge the MF and MB strategies. Alongside, we review hippocampal sequences such as hippocampal replay, theta sequences and temporal coding in the hippocampus, demonstrating how the hippocampus can implement MB learning. We argue that the implementation of such sequences in reinforcement learning agents improves their performance, as demonstrated by both successor representation and continual learning agents. The incorporation of sequences enables these agents to perform not only more complex tasks, but also multiscale temporal processing in a biologically plausible manner. Altogether, we articulate an integrative framework to advance current striatal-focused decision making to better account for multiscale mechanisms underlying various real-world time-related concepts such as the self that cumulates over a person's life course.

### P1-H-30 - Reward-Driven Superstitions in Uncontrollable Environments

Haoyang Lu<sup>1</sup>, Tianyuan Teng<sup>1</sup>, Hang Zhang<sup>1</sup>

<sup>1</sup>Peking University

**Objective:** Superstitions are false beliefs about causality, often related to an illusory control of luck. Although superstitions are ubiquitous and relevant to our everyday lives, previous research has mainly focused on their historical and cultural background, individual differences, and influencing factors. The cognitive processes underlying the formation and persistence of superstitions have been understudied. In this study, we aimed to quantify and investigate the formation of superstitious beliefs in a freely explorable but uncontrollable environment.

**Methods:** We designed a behavioral task with different reward contingencies to track the evolution of superstitious beliefs over time. The task involved three between-participant conditions that differed in the mean and variation of reward probability. We recruited 60 participants for each condition. During the task with 100 trials, participants attempted to produce rewardable keypress patterns that were the combination of four keys by themselves. However, unbeknownst to them, the outcomes were independent of their actions. After the task, participants reported their perceived controllability and reward probability during the experiment. They also completed several questionnaires related to their inclination to have superstitious beliefs in real life.

**Results:** To quantify the evolution of superstitious beliefs, we used multidimensional scaling to map the distances between keypress sequences to a new two-dimensional space. The 340 originally different sequences were then classified into distinct 50 keypress patterns. Similarity analysis of keypress patterns revealed that participants adopted the win-stay, lose-shift strategy ( $F(1,179)=227.13, \hat{1}\cdot g2=.11, p<.001$ ), a heuristic to maximize rewards. Those under high reward rates preferred the different sequences within the same pattern after losing and converged to fewer patterns compared with other conditions ( $F(2,179)=8.89, \hat{1}\cdot g2=.016, p<.001$ ). Although they had no control of the rewards, participants' perception of controllability was higher under high reward rates, but it increased with their perceived, rather than actual, reward rate (for perceived reward rate:  $\hat{1}^2=0.39, t(177)=3.72, p<.001$ ; for actual reward rate:  $\hat{1}^2=-0.51, t(177)=-1.49, p=.38$ ). We also found that the strength of participants' belief in everyday superstitions could predict their reported controllability ( $\hat{1}^2=0.18, t(168)=2.53, p=.012$ ).

**Conclusions:** By quantifying the evolution of human behaviors in uncontrollable environments, we find that the formation of superstitious beliefs may be motivated by reward maximization but is associated with a distorted perception of the reward rate.

### P1-H-31 - Ambiguity and confirmation bias in reward learning

Hayley Dorfman<sup>1</sup>, Rahul Bhui<sup>2</sup>

<sup>1</sup>Harvard University, <sup>2</sup>Massachusetts Institute of Technology

**Objective:** We tend to interpret ambiguous feedback in ways that confirm our preexisting beliefs. We propose that this confirmation bias constitutes an inductive bias that speeds up learning, analogous to missing data imputation. We develop and test this theory using a reinforcement learning (RL) task in which participants are only provided partial information about outcomes, allowing more leeway for subjective interpretation.

**Method** We introduce a computational model of reward learning from feedback that is sometimes ambiguous. Specifically, we consider valence ambiguity, where people are told the exact magnitude but not whether it is positive or negative. The agent imputes these outcomes with an estimated reward and corresponding uncertainty, according to Bayesian reasoning about the outcome's expected value given the partial information combined with their prior beliefs. To test this model, we developed a behavioral task where participants ( $N = 68, \text{MTurk}$ ) chose between two gold mines. After each choice, they saw either a positive (gold), negative (rocks), or ambiguous (dirty sediment) outcome and reported whether they thought the ambiguous outcome was positive or negative.