

## **Phasic dopamine release during reward learning under uncertainty.**

Andrew S. Hart and Paul E. M. Phillips

Program in Neurobiology & Behavior and Departments of Psychiatry & Behavioral Sciences and Pharmacology, University of Washington, Seattle, WA 98195, USA.

Humans and animals are constantly faced with decisions between prospects whose outcomes are risky or probabilistic. It is often the case that the probability and magnitude of potential outcomes are not known a priori but must be learned through repeated interactions with the environment. In well-trained primates, it has been shown that midbrain dopamine neurons fire bursts of action potentials in response to visual conditioned stimuli (CSs) that are followed by juice rewards with some probability. The magnitude of these CS-evoked increases in firing rate scale monotonically with reward probability, as well as reward value, suggesting that they represent expected value (EV) of the reward. In addition, these neurons show sustained firing between CS presentation and reward, which scales with reward variance or uncertainty of outcome. This profile of midbrain neuronal activity is consistent with the hypothesis that dopamine encodes reward prediction error (RPE) signals, which can be used to update a representation of the value associated with the CS. While these responses are well characterized in well-trained animals, little is known about the time course of their acquisition and their relationship to the acquisition of behavioral responses to the CS.

Using fast-scan cyclic voltammetry (FSCV) in rats chronically implanted with carbon fiber electrodes in the Nucleus Accumbens, we recorded dopamine release over the time course of a classical conditioning paradigm in which a light/lever CS is followed by a food pellet reward with some probability (0.25, 0.5, 0.75, or 1). We found that rats trained on this probabilistic reinforcement paradigm developed dopamine responses to the CS over a similar time course, across probabilities, and that the CS responses scaled with probability. The rats also showed dopamine responses to the reward that scaled inversely with probability, suggesting that they represent RPE signals. Sustained dopamine responses between CS onset and reward presentation were also present, but their acquisition was retarded relative to that of CS signals. These results suggest that during classical conditioning a representation of the EV develops before a representation of the uncertainty associated with a CS.

# Memory-based decision making: hippocampally-linked representations underlie behavior in a rewarded choice task

A. M. Bornstein<sup>1</sup> and N.D. Daw<sup>1,2</sup>

<sup>1</sup>Department of Psychology, <sup>2</sup>Center for Neural Science, New York University

Objective: Two distinct neural representations of probabilistic associations are formed during an unrewarded sequential learning task. How are these used when subjects are asked to make rewarded decisions?

Methods: We examined how reaction times (RTs) and - in a second experiment - decisions for money were influenced by associative learning about nonrewarded stimulus-stimulus or response-response contingencies in a sequential image identification task. Over the length of the task, RTs correlated with (uninstructed) image-image transition probabilities. Computational models were fit to the course of learning. Having identified correlates of learning in the absence of reward, we examined how these associations could be combined with rewards to guide decisions. In a second behavioral experiment, we sporadically interrupted the picture sequence with probe choice trials under novel reward contingencies, allowing subjects to draw on their picture sequence learning to obtain reward. Because rewards were specific to the probe, the task precluded decision by standard model-free reinforcement learning (RL) approaches.

Results: RTs were significantly better described by a mixture of two error-driven learning processes with drastically different learning rates. Using fMRI, we identified distinct correlates of each process in hippocampus and ventral striatum. For the reward trials, we leveraged our RT results as signatures each system and identified the model that best explained choices separately from RTs. A single process learning at the predetermined slow rate (the one consistent with hippocampal BOLD) described choices more accurately than the fast (striatal) rate, a freely varying rate fit solely to choices, or any combination of these.

Conclusions: These results suggest that hippocampally-dependent stimulus-stimulus representations are used in reward-based decision making in a task that encourages forward planning of a type not supported by model-free RL. They also demonstrate a novel link between reward-based decision making and hippocampal processes. Finally, they support the notion that model-free RL is insufficient for describing all human choice behavior, and provide a platform for further development of choice models describing more complex decisions requiring memory-guided planning.

Acknowledgements: This study was funded by a McKnight Scholar Award and a NARSAD Young Investigator Award.

# **Time-course of Encoded Expected Utility Revealed by Single Neuron Activity in the Human Amygdala**

Rick L. Jenison<sup>1\*</sup>, Christopher K. Kovach<sup>2</sup> and Martina Chura<sup>3</sup>

<sup>1</sup>Department of Psychology, University of Wisconsin, Madison, WI; <sup>2</sup>Department of Neurosurgery, University of Iowa, Iowa City, IA; <sup>3</sup>Department of Economics, University of Wisconsin, Madison, WI

\*Correspondence at: rjenison@wisc.edu

**Objective:** The time-course of an encoded decision variable leading up to the point of choice cannot be directly observed but can be inferred from neural activity preceding the choice. Naive examination of the correlation between choice and integrated neural responses necessarily assumes a fixed and constant (“boxcar function”) form for the decision variable over the window of integration, and thus relies on an implausible model of the temporal evolution of the signal. In order to reconstruct the time course of value encoding, we replaced the boxcar function with an orthogonal series of polynomials within a conditionally independent binomial-GLM. This approach yields a new semiparametric technique for estimating the evolution of an encoded decision variable based on single neuron activity. We apply the method to human intracranial recordings to show the dynamic evolution of uncertainty-related value signals encoded by single neurons in the amygdala.

**Methods:** We investigated the neural basis of uncertainty coding by recording single neuron activity in the human amygdala while patient-participants chose between the opportunity to gamble or accept a sure win. These studies were carried out in patients undergoing diagnosis and, later, surgical treatment for medically intractable epilepsy. Participants were awake and alert during recording sessions. We recorded from microcontacts implanted directly in the amygdala nuclei. We modeled the neural spiking activity using a generalized linear model (Binomial-GLM) along with a nonlinear stochastic choice model based on revealed choices. The time-course of the hypothesized decision variable was identified by introducing a set of orthogonal Laguerre basis functions as covariates into the GLM that expand the point of choice backward in time. The basis order and the Laguerre pole, which determines the rate of exponential decay, were set to support coverage of the longest single trial duration. The technology was validated using simulated Poisson process spike trains generated from a known time-varying expected utility.

**Results:** A nonmonotonic time-course was commonly observed with a modal region between 500 msec and 1000 msec prior to the revealed choice. The GLM fit of the Laguerre-expanded decision variable improved significantly over that of the boxcar encoding in nearly all cases using the likelihood ratio test.

**Conclusions:** The described polynomial expansion of neurally encoded decision variables in the GLM modeling of spike rate provides two advancements: (1) it gives a principled alternative to the assumption of boxcar-like encoding of decision variables, and (2) it reveals the time-evolution of encoding.

## **Acknowledgements:**

This study was funded by the NSF, grant SES-0820316 and NIH, grant NIAAA R01AA018736 (CRCNS)

## **Appetitive State Toggles a Neuronal Switch for Approach/Avoidance**

Keiko Hirayama and Rhanor Gillette

Neuroscience Program and the Department of Molecular & Integrative Physiology, University of Illinois at Urbana-Champaign, Urbana, IL

Goal-directed neuronal networks may express appetitive state through integrating sensation, internal state and memory, and toggle decision by corollary outputs to competing and effector networks. These relationships are accessible at the level of the networks and their identified neurons for the foraging decisions of the predatory sea-slug *Pleurobranchaea californica*.

Appetitive state is conserved in the isolated CNS, where the frequency of slow spontaneous rhythmic activity in the feeding motor network is a logarithmic relation of donor feeding thresholds. Donor CNS with high feeding thresholds (low readiness to feed) express fictive avoidance turns in response to sensory nerve stimulation; whereas hungry animals (low feeding thresholds) express orienting turns. Enhancing feeding network excitation by depolarizing feeding command neurons also reversibly changes avoidance to orienting. Serotonin added at 5  $\mu\text{M}$  switches avoidance turns to orienting. Serotonin is an intrinsic modulator of both feeding network excitation and appetitive state, and its neuronal content varies with satiation state.

We find that corollary outputs from the feeding network target identified neuron connections of the turn motor network, biasing the symmetry of the turn network from avoidance to orienting in response to sensory input.

Thus, appetitive state is manifest in the excitation state of the feeding motor network, which sets sensory feeding thresholds and biases the turning motor network between approach and avoidance. The simple neuronal mechanisms of appetitive state and approach/avoidance decision in this model system may well be part of the core of more complex economic decision-making, and provide a basis for autonomous cost-benefit decision in artificial-life modeling and robotics.

### Acknowledgements:

Support by NSF IBN 0843621 and NIH R21 DA023445.

## The Nature of Saliency in Strategic Games: Predictions from Visual Neuroscience

Milica Milosavljevic<sup>1\*</sup>, Alec Smith<sup>2\*</sup>, Christof Koch<sup>1</sup>, & Colin Camerer<sup>2</sup>  
<sup>1</sup>CNS and <sup>2</sup>HSS, Caltech, Pasadena, CA, \*Equal Contribution

Correspondence: mmilosav@hss.caltech.edu

**Objective:** In a coordination game players choose objects from a set and prefer to choose objects that other players choose. Since Schelling (1960), the explanation for apparently high levels of coordination has been that some strategies (object choices) are more “salient” than others: they are “psychologically prominent” or “focal”. Many experimental studies have explored strategy choices with indirect reference to saliency, without proposing precise metrics that predict what is salient (and what is likely to be chosen) purely from stimulus properties. We introduce the idea that saliency in games of choice among visually-presented objects is driven by the concept of visual saliency that is well understood in psychology and neuroscience. Visual saliency refers to early, automatic attention directed towards objects in a scene that “pop out” based on low-level visual features such as brightness, color, orientation, etc. We propose that, in some coordination games, players’ strategic behavior is influenced by low-level visual saliency.

**Methods:** Computational neuroscience offers a neurally-plausible algorithm of visual saliency which creates a “saliency map” predicting the bottom-up time course of visual attention (Itti & Koch, 2001; Walther & Koch, 2006). We use this saliency algorithm to design coordination games and predict their outcomes *ex ante*. Each game consisted of a number of strategies (6-35) represented by a variety of object labels across games (e.g., lines, circles, letters, etc.). Subjects were rewarded for choosing the same strategy as an anonymous co-player. We hypothesized that subjects’ choices will be biased towards strategies with salient labels.

**Results:** Data from N=64 subjects show that in 7 newly designed coordination games subjects most frequently chose the label that was *ex ante* deemed to be salient by the saliency algorithm. We also created 3 games in which the most salient object from one game became more or less salient in other games by manipulating its distance and orientation with regard to its neighbors. As predicted, the percentage of subjects who chose the target object decreased as the object became less salient, and increased as it became more salient. Two control games were created that do not follow the psychologically defined rules of visual saliency. In both games, the saliency algorithm failed to predict the outcome of the game.

**Conclusions:** A common neuromorphic saliency algorithm, based on simple visual features such as orientation, color and intensity, predicts high-level strategic choices with far above chance probability. This provides insights into the biological underpinnings of social behavior. An interesting question is what aspects of more complex social games (e.g. hide-and-seek games;) can also be explained using simple neurobiological plausible operations common to all subjects.

## Neural substrates of delay discounting in smokers and nonsmokers

Sarah Tappon,<sup>1</sup> Vanessa B. Wilson<sup>1</sup> and Suzanne H. Mitchell,<sup>1,2\*</sup>

<sup>1</sup>Department of Behavioral Neuroscience, Oregon Health & Science University <sup>2</sup>Department of Psychiatry, Oregon Health & Science University.

\*Correspondence to: [mitchesu@ohsu.edu](mailto:mitchesu@ohsu.edu)

**Objectives:** Smokers prefer smaller but immediate rewards over larger, delayed rewards more than nonsmokers, that is, they discount the value delayed rewards more steeply than nonsmokers (e.g., Mitchell, 1999, *Psychopharm.* 146, 455-464). Also, smokers prefer smaller, slightly delayed rewards over larger, more delayed rewards more than nonsmokers. However, this group difference is reduced when both rewards are delayed. We used fMRI to explore whether the neural substrates of these types of decisions differed for smokers and nonsmokers.

**Methods:** In an event-related design, regular smokers and nonsmokers chose between a small, immediate amount of money and a larger, more delayed reward (single-delay trials), or between a small, delayed reward and a larger, but more delayed rewards (double-delay trials). Half of the choices were designed to be difficult (alternatives were near each subject's point of subjective equality, based on the subject's past performance) and half were designed to be easy. Contrasts were used to identify differences between smokers and nonsmokers under the different decision making conditions (single- versus double-delay trials) at each level of difficulty.

**Results:** Both smokers and nonsmokers showed a robust response in numerous brain regions previously shown to be engaged by discounting tasks. Single-delay trials produced greater activation than double-delay trials across large swaths of the frontal and parietal cortexes, including anterior cingulate, ventromedial prefrontal cortex, and ventral striatum for smokers and nonsmokers. Hard choices were associated with more activity than easy choices in reward-related regions in the midbrain and ventral striatum and frontal regions including dorsolateral prefrontal cortex. Smokers showed more activity than non-smokers in bilateral insula and right dorsolateral prefrontal cortex, demonstrating that group differences in information processing exist even after controlling for subjective choice difficulty.

**Conclusions:** Results suggest similar levels of performance in smokers and nonsmokers were associated with dissimilar patterns of activation, possibly indicating that chronic exposure to nicotine is associated with reduced efficiency of neural processing.

**Acknowledgements:** This study was funded by R03 DA024195.

## **Decision Making and the Brain: Decision Strategies In the Psychopathic Brain**

D. T. Wargo<sup>1</sup> and G. L. Spencer<sup>2</sup>

<sup>1</sup>Economics Dept., Temple Univ. (docwargo@temple.edu)

<sup>2</sup>Human Resource Mngt. Dept., Temple

Objective: While there has been significant research on psychopaths in institutionalized settings, there is a paucity of studies of non-institutionalized psychopaths, who make up 1% of the general population and 3.5% of business employees. Furthermore, study of the behavior and decision-making strategies of non-institutionalized psychopaths in organizations is virtually non-existent, due to privacy, logistical and bureaucratic roadblocks. We use well-validated psychopathy scales to identify individuals with psychopathic attributes and more importantly observe their decision-making and behavior in real world, long-term, small group (N=10) interactions that require cooperation to achieve real-world payoffs (= grades).

Methods! The participants in the study are 182 students (all freshmen, M= 103, F= 79) in 2 sections of an undergrad Human Resources Management class at the business school of a major research university. The psychopathic questionnaire (titled ‘Success Strategy Assessment’) was administered through a web-based format on Blackboard. The authors are specialized in neuroeconomics, behavioral economics, organizational dynamics and conflict management. We analyzed the ‘free riding’ and team disruption caused by these same individuals in team activities that required cooperation in significant, graded, semester-long projects. Six extensive peer evaluation questionnaires assessing team cohesion, individual contribution and conflict were administered, along with extensive observation by one of the co-authors.

Results: The results are fascinating and contrary to expectations. A surprising number of individuals (both male and female) endorsed psychopathic attributes that were in the second standard deviation above the mean. Further, when we correlated the total psychopathic attributes endorsed in each group with the final course grade, we found that the psychopathic attributes were additive. That is, the higher the total psychopathic score in the group, the lower was the group’s course grade!

Amazingly, the strategies used by the psychopaths in the groups were quite varied. Some were simply free riders. Some hijacked the entire group and compromised the peer reports. Others attempted to manipulate the professor.

Conclusions: The influence of individuals with psychopathic attributes is magnified by group dynamics. This has been demonstrated in research showing that a large portion of the general population are ‘reciprocators’ who behave ethically in an ethical culture but much less so when they observe other individuals “getting away with it”.

## The neural correlates of primary and secondary costs in economic decision-making

Hilke Plassmann<sup>1,2\*</sup>, Nina Mazar<sup>3</sup>, and Antonio Rangel<sup>4</sup>

<sup>1</sup>INSEAD, <sup>2</sup>Ecole Normale Supérieure, <sup>3</sup>University of Toronto, <sup>4</sup>California Institute of Technology

\*Correspondence at: hilke.plassmann@insead.edu

**Background:** A seminal account proposed for the hedonics underlying consumption decisions distinguishes between two opposing hedonic factors: rewarding factors like the pleasure derived from consumption and aversive factors like the costs of the consumption often referred to as “pain of paying” (Prelec & Loewenstein 1998). Previous research in neuroeconomics has investigated how the positive hedonics in form of goal values at time of choice are represented in the brain (Plassmann et al. 2007) and whether these representations differ for different modalities such as primary or secondary rewards (Chib et al. 2009). However, despite its importance for economic decision-making, little is known about how the human brain computes costs during decision-making. In particular, it is unknown if different types of costs (e.g., monetary costs vs. physical pain costs) are processed by the same brain networks. This is the central question of this paper. In particular, we investigated whether that representation differs between abstract (secondary) costs (e.g. paying money) or somatosensory (primary) costs (e.g. tolerating electric shocks) that are matched in economic value.

**Background:** In this paper we investigated the neural basis of cost computations by scanning hungry subjects’ brains (N=21, aged 18-35, mean 23.65 years) while making 280 purchasing decisions. The fMRI task consisted of two different trial types: Trials in which subjects could decide to buy 40 food items at four different monetary prices (\$0, \$0.50, \$1.00, \$1.50) (= 160 ‘\$ trials’) and trials in which they could decide to buy the same 40 food items for tolerating pain (electric shock) at three different pain intensities (=120 ‘V trials’) that were matched with the three different non-zero monetary prices.

**Results:** For the behavioral data analysis, we created dummy variables for WTP and purchasing prices in money and physical pain trials and entered them into a mixed effects logistic regression analysis. We tested differences in the regression coefficients between money and physical trials for WTP and price predictors and found significant differences (both  $p < .001$ ). For the fMRI data analysis we estimated a hierarchical mixed effect GLM to investigate differences and overlaps for brain areas that correlated with the size of monetary and physical pain costs. We found that the bilateral insula and the thalamus, regions involved in pain processing, correlated positively with the size of physical pain prices, but not with the size of monetary prices ( $p < .001$ , uncorr.). A conjunction analysis revealed that no overlapping areas can be found ( $p < .001$ , uncorr.).

**Conclusion:** Taken together, these results show that people react differently to “monetary” and “pain costs” on a brain, but not a behavioral. Our fMRI results suggest that paying with money might trigger very different emotional processes than those involved with more “physical” forms of costs. These results call into question “pain of paying” theories recently suggested in the neuroeconomic literature at least for everyday consumption decisions. Our results have important implications for disadvantageous decision-making such as overspending and shopping addiction.

## **An addiction to stuff? An fMRI study to determine if common neural circuits underlie decisions about cotton balls and crack**

John M. Wang\*<sup>1</sup>, Rachael D. Seidler<sup>2,3</sup>, Julie L. Hall<sup>2</sup>, and Stephanie D. Preston<sup>2</sup>

<sup>1</sup>University of California, Davis, Center for Mind and Brain

<sup>2</sup>University of Michigan, Department of Psychology

<sup>3</sup>University of Michigan, School of Kinesiology

\*Correspondence: wangjm1985@gmail.com

**Objective:** People's decisions about goods extend well beyond new purchases. Every pencil, book, t-shirt, or gadget acquired entails a chain of decisions regarding where to put it, for how long, and what to do with it when it is no longer valued. Despite a mainstream interest in hoarding, little is known about the neural bases of acquiring and discarding common household goods. One could assume that such decisions are subserved by the mesolimbocortical system; however, most decision-making studies only investigate the acquisition phase and entail a monetary exchange. In contrast, we know that price is not germane to the selection and retention of common goods and work on compulsive hoarding has not found a role for the nucleus accumbens (NAcc). To directly study this issue, we contrasted decisions to acquire versus discard common goods, while maximizing personal preference vs. monetary profit.

**Methods:** Twenty participants (11 males) completed a block-design fMRI study using a forced-choice procedure under four nested frames. We contrasted acquisition to discarding and personal to monetary decisions (using response time as a nuisance covariate), and regressed BOLD signal responses in the regions of interest with trait measures of hoarding.

**Results:** The orbital frontal cortex (OFC) was involved across frames, but all other regions were significantly affected by the frame, which also changed qualitative object preferences. Monetary decisions recruited executive regions, and biased choice towards valuable (but less desired) items. Personal decisions recruited midline, affective regions and caused subjects to emphasize inexpensive, immediate rewards (candy, coins). Acquisition augmented OFC activity and biased subjects towards shiny, metal objects while discarding recruited the anterior cingulate and insula, and biased subjects towards utilitarian, "should" items. In contrast to neuroeconomic studies, acquiring such goods did not engage the NAcc, which instead was only activated during personal acquisition with the degree of trait hoarding tendencies.

**Conclusions:** Decisions about everyday goods may necessarily involve the OFC, but otherwise represent a complex interplay between the frame, the item, and the decider. Perhaps for acquisitive individuals even mundane goods acquire the incentive salience of more hedonic items (money, food, luxury items). Research on material goods is critical to understand a ubiquitous human decision that is critical to our economy, our quality of life, and the environment.

### **Acknowledgements:**

This study was funded by a grant from the University of Michigan.

# The Time Course of Value Computations at the Time of Decision-Making

Alison Harris\*, Ralph Adolphs, Colin C. Camerer, Antonio Rangel

California Institute of Technology

\*Correspondence at: [amharris@caltech.edu](mailto:amharris@caltech.edu)

Objective: In the decision-making literature, it is commonly held that the brain makes simple choices by assigning values to the stimuli under consideration, and then comparing them to select the best one. Growing evidence from electrophysiology and neuroimaging has implicated specific brain regions in value signal computations, most notably ventromedial prefrontal cortex (vmPFC) and anterior cingulate cortex (ACC). Yet the inputs to this computation, as well as its time course of activation, are unknown.

Methods: We examined the time course of value signal computations by measuring event-related potentials (ERP) while hungry subjects performed a simple choice task involving a variety of appetitive and aversive foods.

Results: Linear ERP responses consistent with value coding were observed as early as 150-250 ms after stimulus onset, with a highly significant linear response between 400 and 550 ms at central and frontotemporal sensors. We also found later frontal activity from 700-800 ms, coinciding with the average median response time (710 ms). This late response varied with differences in reaction time, showing no significant linear effect when reaction times were fast. Finally, a distributed source reconstruction localized early activity to areas of medial temporal lobe, insula, and superior frontal gyrus, with vmPFC sources observed beginning in the 400-550 ms range. Over the time window from 400 ms until response, localized source activity appeared to spread from posterior to anterior vmPFC, and from vmPFC to ACC, with more pronounced effects for trials with longer RT. Similarly, a response-locked analysis showed significant localization to vmPFC sources from approximately -240 to -160 ms, and ACC sources from about -160 to -80 ms, before response onset.

Conclusions: Our results suggest that vmPFC contributions to value signal processing reflect integration of value-related information from the medial temporal lobe, and occur relatively late in the decision process. These data provide preliminary evidence regarding the time course of value signal computations in the human brain.

## **The Neuroeconomics of Nicotine Dependence: Understanding Intertemporal Choice in Smokers using fMRI**

Michael Amlung, MS<sup>1</sup>; James MacKillop, PhD<sup>1</sup>; Lawrence Sweet, PhD<sup>2</sup>; Lauren Wier, MPH<sup>2</sup>; Sean David, MD, SM, DPhil<sup>3</sup>; Beth Jerskey, PhD<sup>2</sup>; Lara Ray, PhD<sup>4</sup>; James Murphy, PhD<sup>5</sup>; Ronald Cohen PhD<sup>2</sup>; Warren Bickel, PhD<sup>6</sup>

**Objective:** Behavioral studies have consistently found significantly more impulsive intertemporal preferences in nicotine dependent individuals compared to healthy controls. These studies use delay discounting paradigms, assessing how much a reward loses value based on its delay in time. Studies also reveal a “commodity effect”, reflecting greater impulsivity for choices involving cigarette rewards. Several recent studies have examined the neuroanatomical basis for delay discounting in healthy adults using fMRI, but none have examined smokers or commodity effects, which are the foci of the current study.

**Methods:** The study used a within-subjects design and enrolled 15 right-handed nicotine dependent adults with no history of brain injury or MRI contraindications. Participants underwent a rapid event-related design during which they responded to 108 choices for smaller and larger commodities (half money and half the equivalent number of cigarettes) available immediately or after a delay period. Stimuli comprised 72 experimental items (smaller immediate vs. larger delayed option) and 36 control items (smaller immediate vs larger immediate). The study used a Siemens 3T TIM Trio MRI. Scanning parameters were axial orientation, TR 2500 ms, TE 28 ms, FOV = 192<sup>2</sup> mm, matrix = 64<sup>2</sup>, slice thickness = 3 mm.

**Results:** Responses to the experimental items were classified as “impulsive” (preference for smaller immediate reward) or “restrained” (preference for larger delayed reward). Primary analyses used a disjunction mask to identify regions associated with significant activity in any of the three categories of choices (impulsive, restrained, control), identifying 23 empirical regions of interest (ROIs). These were then examined across categories to identify regions that discriminated across choice type, revealing 13 significant or trend-level effects. To reduce Type I error, a modified Bonferroni correction was applied (Hochberg step down procedure), leaving five ROIs, including the middle and medial frontal gyri, cerebellum, and precuneus. Commodity effect analyses used voxelwise paired t-tests irrespective of choice category for maximum power, revealing significant differences in the medial, middle, and inferior frontal gyri, cerebellum, and the inferior parietal lobule.

**Conclusions:** The results are consistent with prior studies of healthy adults, suggesting common neuroanatomical processing across samples. Commodity effects were evident and suggest biases toward tobacco cues in prefrontal regions associated with decision-making and cognitive control. Future comparative studies and further examination of commodity effects are warranted.

**Contact:** James MacKillop, PhD; email: [jmackill@uga.edu](mailto:jmackill@uga.edu); telephone: 706.542.3033

---

<sup>1</sup> University of Georgia

<sup>2</sup> Brown University

<sup>3</sup> Stanford University

<sup>4</sup> University of California, Los Angeles

<sup>5</sup> University of Memphis

<sup>6</sup> University of Arkansas for Medical Sciences

## When you keep changing your mind: The neural basis of preference reversals

Joseph W. Kable\* and Jessica Stump

Department of Psychology, University of Pennsylvania

\*Correspondence at: kable@psych.upenn.edu

**Objective:** Systematic inconsistencies in people's decisions provide a central challenge to rational choice theories. A classic example is the "preference reversal phenomenon" (Grether & Plott, 1979; Tversky, Slovic & Kahneman, 1990), when people choose one gamble over another, but then bid a higher price for the non-chosen gamble. Some explanations of preference reversals emphasize task-dependent changes in valuation—e.g., that the weight given to payoff versus probability information changes between choice and bidding. Other explanations focus on differences in the decision process between tasks—e.g., that people overbid on certain gambles because of an anchoring heuristic. We gathered neuroimaging data to inform the debate about the cause of preference reversals.

**Methods:** Preference reversals were elicited using a risky gamble paradigm in 24 subjects undergoing fMRI. In alternating scans, subjects made choices between gamble pairs or bid on single gambles. Both tasks were administered in an incentive compatible manner. Each choice involved two gambles of equal expected value, shown sequentially: one had a high probability of winning a small amount of money (termed the *P-bet*, e.g., 84% chance of \$20), and another had a low probability of winning a greater amount (termed the *\$-bet*, e.g., 24% chance of \$70). The same gambles were used in the bidding task.

**Results:** Consistent with previous results in this paradigm, and with general findings of risk aversion, subjects chose the *P-bet* 65% of the time. However, when bidding on the gambles separately, subjects assigned a higher dollar value to the *\$-bet* 67% of the time. Preliminary analyses identified widespread regions that showed greater activity for the preferred gamble category in each task (*P-bets* in choices, *\$-bets* in bids). In medial and ventromedial prefrontal cortex and superior frontal sulcus, this effect was greater for those gamble pairs where subjects made a reversal compared to those pairs where they remained consistent across tasks. The size of this effect was also correlated across subjects with the number of preference reversals made. Furthermore, these effects were primarily driven by modulations in activity specific to the bidding task.

**Conclusions:** These results demonstrate that preference reversals are accompanied by changes in neural activity in medial prefrontal regions previously shown to encode subjective value (Kable & Glimcher, 2009), consistent with an explanation based on task-dependent changes in valuation. The specificity of the neural effects to bids further suggests that these changes could arise because of a cognitive process differentially engaged during bidding.

# The value computations in vmPFC and the striatum are guided by visual attention

Seung-Lark Lim<sup>1</sup>, John P. O'Doherty<sup>2</sup> and Antonio Rangel<sup>1\*</sup>

1. California Institute of Technology, 2. Trinity College, Dublin, Ireland

\* Correspondence at [rangel@hss.caltech.edu](mailto:rangel@hss.caltech.edu)

When facing a choice between two or more stimuli, individuals typically evaluate the options by sequentially deploying their gaze to the different options, often by going back and forth. This behavioral observation suggests that visual attention might play a critical role in the value computations that are carried out in areas such as ventromedial prefrontal cortex (vmPFC) at the time of choice.

We present the results of a study design to test the hypothesis that visual attention selectively modulates stimulus value signals in ventromedial prefrontal cortex (vmPFC) and ventral striatum (vStr) at the time of decision-making. During 3-day fMRI sessions, participants (N=20) performed a real binary choice task. Every trial they were presented with pictures of two food items and were asked to alternate the focus of visual attention (eye-fixation) between the two items. The duration of each eye-fixation randomly varied from 1 to 4 sec. The vmPFC and vStr ROIs were identified by using a parametric regressor of stimulus values acquired during a separate liking-rating task.

We found that the value computations in the vmPFC and vStr were strongly modulated by the visual attention manipulation: activity correlated with the parametric regressor ( $V^L - V^R$ ) when participants fixated the left item, while those ROIs were negatively correlated with this same regressor when they looked to the right item. Furthermore, a PPI analysis showed that the vmPFC exhibited significant functional connectivity with the Left STG/MTG during the left-fixation condition, and with R STG/MTG during the right-fixation condition.

Our results suggest that visual attention plays a critical role in the value computations in the vmPFC and vStr at the time of choice: the value signals in these areas seem to encode a “value of attended minus value of unattended” signal that could be useful in guiding choice.

Acknowledgement: This study was funded by the Moore Foundation

## Genes, Economics, and Happiness

Jan-Emmanuel De Neve,<sup>1\*</sup> James H. Fowler,<sup>2</sup> Bruno S. Frey<sup>3</sup>, and Nicholas A. Christakis<sup>4</sup>

<sup>1</sup>London School of Economics; <sup>2</sup>University of California, San Diego; <sup>3</sup>University of Zurich; <sup>4</sup>Harvard Medical School.

\*Correspondence: [j.c.de-neve@lse.ac.uk](mailto:j.c.de-neve@lse.ac.uk)

Objective: To show the importance of genetic variation in individual subjective well-being and test a specific genetic polymorphism in the serotonin transporter gene for its neurological influence on happiness.

Methods: Using data from the National Longitudinal Study of Adolescent Health (Add Health), we first employ **twin design** methods that compare the behavior of 490 same-sex twin pairs (271MZ, 219DZ) to gauge the part of the variance in happiness (as measured by life satisfaction) that can be attributed to additive genetics (A), common environment (C), and unique experience (E). Next, a **candidate gene association** study (N=2,574) tests a functional polymorphism in the promotor region of the serotonin transporter gene (5-HTTLPR) given prior association with mood regulation, mental health, and selective processing of positive and negative emotional stimuli.

Results: First, the ACE model puts heritability ( $h^2$ ) of life satisfaction at approx. 33% [Women, 26%, Men 39%]. While important, it suggests the need to revise downwards previous estimates of  $h^2 = 50\%$ . Preliminary longitudinal modeling suggests that the  $h^2$  of happiness increases with age. Second, we find that **individually with the transcriptionally more efficient version of the serotonin transporter gene (5-HTTLPR) are significantly more likely to report higher levels of life satisfaction** ( $p=0.005$ ). Having one or two alleles of the more efficient type raises the average likelihood of being very satisfied with one's life by 8.5% and 17.3%, respectively. We replicate this association on an independent sample (Framingham Heart Study) using the linked rs2020933 SNP. We seek to further replicate this association.

Conclusions: Generally, the importance of integrating the study of genetic variation in order to advance our understanding of the biological underpinnings of individual behavior. Specifically, these results may help explain the important genetic component of the individual baseline levels of happiness.

## OFC Value Neurons Do Not have Spatial Tuning

Lauren Grattan,<sup>1</sup> and Paul W. Glimcher<sup>2</sup>

<sup>1</sup>Center for Neural Science, New York University; <sup>2</sup>Center for Neuroeconomics, New York University

Objective: Only a limited amount of information is available about the value representation encoded by single neurons in orbitofrontal cortex. There is overwhelming evidence that neurons in this area have an ordered gustatory response to flavors when monkeys must make choices (Tremblay & Schultz, 1999; Kobayashi S, et al., 2010; Padoa-Schioppa & Assad JA, 2008). However, no group has rigorously tested the spatial tuning properties of these neurons. The degree of spatial tuning for movement metrics is of significant importance to the widely discussed hypothesis that these neurons may support choice in an action-independent manner. Here we report a formal test of the hypothesis that OFC neurons have no spatial tuning for stimulus location or action metrics while retaining strong gustatory tuning in a traditional single target visual-saccadic task.

Methods: Macaque monkeys were trained to perform a visually-guided saccade task with stimuli placed at 1 of 81 locations across the visual field spanning the central 32 deg of visual space. Targets were presented in 1 of three colors and each color was consistently associated with one of three juice flavors (apple, grape, orange). To test the spatial tuning hypothesis we fit the measured response fields with three models: i) A global mean firing rate with no spatial structure. ii) A 2-dimensional plane that could be pitched in any direction. iii) A 2-dimensional Gaussian. We used the Akaike Information Criterion to compare the model types.

Results: We separately analyzed all neurons for evidence of spatial tuning. An average of 548 trials spanning the 81 targets crossed with 3 flavors were examined for the most densely sampled group. The Akaike Information Criterion indicates that the mean firing rate model (the model indicating that there is no spatial tuning – not even contralateral preference) is significantly preferred to all other models for more than 90% of the neurons we examined.

Conclusions: This analysis indicates that OFC neurons show no spatial tuning under these conditions; there does not even appear to be a bias for contralateral target representations. These same neurons do, however, maintain a value-related gustatory response even in a task that does not require a choice between targets yielding different rewards. We conclude that the representation encoded by these neurons is appropriate for the participation in a goods-based choice process that could be coupled to previously identified action-based choice mechanisms in the fronto-parietal choice network.

Acknowledgments: This study was funded the NIE grant R01EY010536.

# **BOLD Response During Passive Viewing of Stimuli Predicts Subsequent Economic Choice**

Alec Smith,<sup>1\*</sup> Doug Bernheim,<sup>2</sup> Colin Camerer,<sup>1</sup> and Antonio Rangel<sup>1</sup>

<sup>1</sup>Division of Humanities and Social Sciences, California Institute of Technology;

<sup>2</sup>Department of Economics, Stanford University

\*Correspondence at: [acs@hss.caltech.edu](mailto:acs@hss.caltech.edu)

Objective: We studied whether brain activity during passive viewing of food images contains information that is predictive of later choices. The study has two objectives. First, does the brain compute value automatically, or are values only computed at the time of choice? We also ask a methodological question: can whole-brain BOLD signals acquired during passive viewing of choice items be used to predict subsequent choices?

Methods: Seventeen hungry subjects passively viewed images of foods during functional imaging. After scanning, the subjects made choices from pairs of the foods shown in the scanner. We extracted the BOLD signals corresponding to each food and used these signals together with subject choices to train a penalized logistic classifier to predict out-of-sample choices from BOLD signals. We used an elastic net (Zou and Hastie 2005) penalty. This classifier both selects predictive voxels, and also allows the contributions of correlated voxels to be averaged. Therefore, it is particularly appropriate for use with fMRI data.

Results: The classification procedure correctly predicted greater than 60% of subjects' choices (between subject average). This average prediction rate is significantly better than chance. On an individual level, prediction rates for a majority of subjects were significantly greater than chance at a 5% threshold.

Conclusions: These results suggest that neural activity during passive viewing contains information that is predictive of future choices. The results provide evidence that the brain automatically encodes components of value, and that these components are both observable and predictive of choices.

Acknowledgements: The National Science Foundation funded this study.

## Frontal Asymmetry and Delay Discounting

Jacob B. Hirsh

Kellogg School of Management, Northwestern University

Correspondence at: [j-hirsh@kellogg.northwestern.edu](mailto:j-hirsh@kellogg.northwestern.edu)

**Objective:** fMRI research indicates that activity in the midbrain reward circuits is associated with preferences for smaller but immediate rewards in a discounting task. There has been comparatively less research, however, on how EEG indicators of reward processing relate to such choices. The EEG frontal asymmetry literature suggests that relatively greater left frontal cortical activity is associated with heightened approach motivation and a stronger engagement with potential rewards. The current study examined whether increased left frontal activity during a delay discounting task would be associated with higher discounting rates.

**Methods:** Forty-eight undergraduate students participated in the study. EEG was recorded from 32 Ag/AgCl electrodes as participants completed a computerized version of the Monetary Choice Questionnaire, which involves a number of choices between monetary amounts of various sizes and delays. Asymmetry scores were calculated as the difference in log transformed power in the alpha frequency band (8-13Hz) in homologous electrode locations across hemispheres. These scores were then correlated with delay discounting rates, as estimated with logistic regression. A measure of cognitive ability was included to examine its relationship with discounting rates and frontal asymmetry.

**Results:** Relatively greater left frontal cortical activity, which has previously been associated with heightened approach motivation and the behavioral activation system, was significantly related to higher discounting rates. Individuals displaying increased left frontal activity during the discounting task were more likely to choose the smaller, but immediate rewards. As in previous work, cognitive ability was related to lower discounting rates. However, regression analyses indicated that cognitive ability and frontal asymmetry were independent predictors of delay discounting.

**Conclusions:** These results suggest that heightened approach motivation (as reflected in greater left frontal activity) is associated with higher levels of delay discounting, consistent with previous individual differences research exploring the roles of extraversion and positive affect. Additionally, these effects were independent of cognitive ability, further emphasizing the role of motivational factors. Finally, the study indicates that EEG may be a useful technique in the investigation of delay discounting processes.

# **The Effect of Short-term Affective Modulation on Reward Prediction Error Signal: A Study of Feedback-related Negativity**

Chun-Yu Chen<sup>1</sup>, Nai-Shing Yen<sup>12\*</sup> and Ruey-Ming Liao<sup>123</sup>

<sup>1</sup> Department of Physiology, <sup>2</sup> Research Center for Mind Brain and learning, <sup>3</sup> Institute of Neuroscience, National Cheng-Chi University, Taipei, Taiwan.

\*Correspondence at: [nsy@nccu.edu.tw](mailto:nsy@nccu.edu.tw)

**Objective:** Previous study has shown the effect of short-term affective modulation on ERN (error-related negativity) induced by Flanker task. However, whether short-term affective modulation can affect on FRN (feedback-related negativity) remains unclear. The present study applied pictures of IAPS and tested the effect of short-term affective modulation on FRN.

**Methods:** Eighteen adult subjects participated in the study. The subject received emotional stimuli via pictures of IAPS presented prior to a two-alternative choice based on a non-reinforcement learning task. In which, feedback would be delivered after every choice made. Three blocks, each of 120 trials, were designed with 75%, 50% and 25%, respectively, as the rewarded feedback. Via ERP measurements, LPP (late positive potential) and FRN were extracted during the presentation of pictures and feedbacks, respectively.

**Results:** As expected, we found that FRN was significantly affected by short-term affective stimuli. However, while FRN was significantly reduced by the positive pictures as compared to the neutral ones, there was no significant difference when comparing the negative and neutral ones. LPP examination showed that the largest amplitude elicited by presentation of positive pictures, which indicated a significant arousal involved. FRN amplitudes were found significantly different over the three conditions of reward feedback. Larger FRN amplitude was found when the feedbacks were unexpected to the subject, whereas smaller FRN amplitude appeared when the feedbacks were expected. There was no interaction between different levels of expectancy and pictures valence.

**Conclusions:** These results suggest that FRN can be modulated by short-term affective stimulation. The subjective arousal level toward affective stimuli might be involved in such a modulatory process.

## Preferences During Repeated Visual Probabilistic Choice

Julia Trommershaeuser<sup>1\*</sup>, Elke U. Weber<sup>2</sup>, Eric J. Johnson<sup>3</sup>, and Paul W. Glimcher<sup>1</sup>

<sup>1</sup>:Center for Neural Science, Dept. of Psychology and Center for Neuroeconomics, New York University

<sup>2</sup>:Department of Psychology, Columbia University

<sup>3</sup>:Columbia Business School, Columbia University

\*Correspondence at: [julia.trommershaeuser@nyu.edu](mailto:julia.trommershaeuser@nyu.edu)

**Objective:** Humans often make demonstrably sub-optimal choices. One of the most frequently observed errors lies in the representation of probabilities. Here we tested the hypothesis, that a visual representation of probability that does not make use of numerical symbols removes the distortions that plague traditional studies of decision-making.

**Methods:** We compared human performance in a repeated ‘visual choice task’ with performance in a traditional ‘numeric lottery choice task’; 8 subjects chose between 174 pairs of gambles that differed with respect to probability and assigned outcomes. Probability information was either coded visually by the density of small gray dots displayed on a white background patch or numerically. In both conditions, information about the magnitude of each possible outcome was presented numerically (amount stated in \$; outcomes > \$0) and through a linear change in the contrast of the dots or the numerals indicating probability. The experiment started with 44 training trials in the visual, and 44 training trials in the numeric choice task. During the visual training period, a scale on the bottom of the screen presented the correspondence between probability and dot density. Subjects chose between the gambles using key presses. Following each choice, the selected gamble was resolved. At the end of the experiment, one of these trials was selected at random for payment to the subject.

**Results:** We fitted the choice data from all trials, individually for each subject, using a single parameter power utility function, and a one-parameter Prelec probability weighting function with a fixed inflection point of  $1/e$  (and a logit error term). Five of the 8 participants showed reduced probability biases for visually compared to numerically coded probability information. Median response times ranged from 0.76 s to 3 s and did not differ significantly for visual and numeric choices. We tested 2 more participants in a modified version of the experiment, in which visual and numeric probability information was interleaved on a trial-by-trial basis. As before, visual representation reduced the biases compared to numerically coded probability information.

**Conclusions:** Our results suggest that visual presentation of probability information and repeated execution of the same decision problem facilitate efficient choice strategies that maximize expected utility. We speculate that distinct neural circuits mediate these two kinds of decision-making; an evolutionarily older and more efficient system that is involved in sensory perception and response generation and a more recent and less efficient symbolic system that is accessed by semantic tasks.

### Acknowledgements:

Grant 1R01AG033406 (NIH/NIA); JT is funded by the Deutsche Forschungsgemeinschaft (through the Heisenberg Programme, grant TR 528/2-1).

# Monetary Reward and Intrinsic Motivation: Neural Basis of Motivation Crowding-Out Effect

K. Matsumoto,<sup>1\*</sup> K. Murayama,<sup>2</sup> M. Matsumoto,<sup>1,3</sup> and K. Izuma.<sup>1,3</sup>

<sup>1</sup>Brain Science Institute, Tamagawa University; <sup>2</sup>Department of Psychology, University of Munich; Japan Society for the Promotion of Science.

\*Correspondence at: matsumot@lab.tamagawa.ac.jp

Opposing the widespread belief that people are positively motivated by reward incentives, some studies have shown that performance-based extrinsic reward can actually undermine a person's intrinsic motivation to interesting tasks. This "undermining effect" (also called "motivation crowding-out effect" in the field of economics) not only has practical implications given the burgeoning of performance-based incentive systems in contemporary society, but also presents theoretical challenge for economic and reinforcement learning theories, which assume that monetary incentives monotonically increases motivation. Despite the practical and theoretical importance of this provocative phenomenon, however, little is known about its neural basis. Herein we induced the behavioral undermining effect utilizing a newly-developed task and tracked its neural correlates using functional magnetic resonance imaging (fMRI). Our results demonstrated that performance-based monetary reward indeed undermined intrinsic motivation assessed by the number of voluntary engagement in the task after the reward provision, and that the activity in the anterior striatum and the prefrontal areas decreased along with this behavioral undermining effect. These findings suggest that the cortico-basal ganglia valuation system underlies the undermining effect through the integration of extrinsic reward value and intrinsic task value. We conclude by urging researchers to incorporate the concepts of intrinsic motivation into neuroscience, as most neuroscience research on motivation to date has been confined to extrinsic rewards such as food or money.

## Acknowledgements:

This study was supported by a Grant-in-Aid for Scientific Research C#21530773 to K. Matsumoto and Tamagawa University Global Center of Excellence grant from the Ministry of Education, Culture, Sports, Science and Technology Japan.

# Discounting Time and Probability by Perception of Reward

William H. Alexander,\* and Joshua W. Brown

Department of Brain and Psychological Sciences, Indiana University, Bloomington.

\*Correspondence at: [wialexan@indiana.edu](mailto:wialexan@indiana.edu)

Objective: A common finding in studies of temporal and probability discounting is that the rate at which future or uncertain rewards are discounted depends on the magnitude of the delayed or uncertain reward. In temporal discounting, large rewards are discounted at a lower rate than small rewards (the magnitude effect), while for probability discounting, larger rewards are discounted at a greater rate (the reverse magnitude effect). Together, the magnitude and reverse magnitude effects have been cited as evidence that separate processes underlie the two forms of discounting. We describe a new model of decision-making, the Recursive Hyperbolic model (RHM), which suggests a common mechanism for both temporal and probability discounting.

Methods: The RHM was fit to published data (*I*) in which human participants indicated preferences for choices involving delayed or probabilistic rewards with two different reward levels. Additional simulations investigated the effect of parameter manipulation on discounting behavior.

Results: Parameters for the RHM were found which accurately captured both the magnitude and reverse magnitude effects observed in human choice. Notably, both temporal and probability discounting were described by a single parameter set. In contrast, previous models of discounting behavior (e.g., hyperbolic discounting) require a separate set of parameters for each reward level. Further simulations suggest that manipulation of a single parameter in the RHM model related to perception of reward produces a pattern of change in discounting behavior consistent with individual differences related to impulsivity. We relate these results to evidence which suggests that dysfunction of the serotonergic system may partially contribute to impulsive behavior.

Conclusions: We present a new model of decision making which suggests a common mechanism, perception of reward, underlies both probability and temporal discounting. The model provides a convenient theoretical framework for studying rewards which are both delayed and uncertain, and provides a novel interpretation of the effects of impulsivity on choice behavior. Furthermore, the existence of a single mechanism subserving both delay and probability discounting suggests both kinds of decision making may share a common neural substrate, and may provide critical insight into the function of neuromodulatory systems implicated in impulsive behavior (e.g., serotonin).

## Acknowledgements:

Supported by AFOSR FA9550-07-1-0454, A NARSAD Young Investigator Award to JWB, the Sidney R. Baer, Jr. Foundation, and R01 DA026457.

## References:

1. J. Myerson, L. Green, J. Scott Hanson, D. D. Holt, S. J. Estle, *Journal of Economic Psychology* **24**, 619 (2003).

# Normalized value coding underlies irrational choice behavior

Kenway Louie and Paul Glimcher

Center for Neural Science, New York University, New York, NY

The neural circuits underlying decision-making must represent the values of the available choice options. In the monkey lateral intraparietal area, neural activity is strongly modulated by the value of specific saccades. In recent neurophysiological experiments, we have demonstrated that this value representation is not absolute: neurons code the value of saccades to the response field relative to the values of all available saccade choices. This value normalization is well-described by a divisive normalization model that also characterizes nonlinear phenomena such as gain control and cross-orientation suppression in visual cortex.

Does this normalized value representation affect behavior? We explore here the predictions of the normalization model and compare them to observed choice behavior. Computational simulation of the choice process indicates that normalized value coding produces a specific violation of rational choice theory: preference between two high-valued options appears to be a function of a third, low-valued irrelevant alternative. This effect depends crucially on cortical neuron response variability: as the total value of available options increases, the separation between the distributions of firing rates representing two differently valued options decreases; if variance does not decrease appropriately, the options will be increasingly difficult to distinguish.

We examined the predictions of the divisive normalization model in both human and non-human primates. We trained two monkeys to choose between three differently valued stimuli (A, B, and C). Stimulus locations and reward associations were fixed within a block; across blocks, the values of the target options (A and B) were varied to quantify how choice varied as a function of value difference. The distractor option (C) provided one of two possible values, which was always lower than any possible target option value. We find that choice is context-dependent: the relative preference between two high-valued options depends on the value of the third option. At low distractor value, monkeys are more likely to correctly choose the better target option; at high distractor value, choice behavior becomes more stochastic. Importantly, this effect is equivalent to a change in the slope of the logistic choice function, an effect predicted by the combination of normalized value and cortical variability. Consistent with a localized cortical mechanism, this context dependence is spatially modulated, exhibiting a stronger effect when the distractor and target options appear in the same visual hemifield. Furthermore, preliminary studies examining human choice behavior in an analogous task show similar results. We conclude that the normalized representation of value in choice circuits is observable at the behavioral level, and may play a role in real-world examples of context-dependent choice.

## Acknowledgements

Funding for this project has been provided by the NIMH and NEI.

# Individual Variance across Probability Discounting, Reversal Learning, and Working Memory in Rats

L.K. Graham<sup>1</sup>, J.K. Chan<sup>2</sup>, D.C. Castro<sup>1</sup>, & J.J. Kim<sup>1,2\*</sup>.

<sup>1</sup>Department of Psychology; <sup>2</sup>Program in Neurobiology & Behavior, University of Washington, Seattle, WA 98195.

\*Correspondence at: jeansokk@u.washington.edu

**Objective:** Decision making is a complex process that requires skill in several cognitive realms, including learning, memory, cognitive flexibility, value detection and comparison, and action selection. Although many decision-making studies address other forms of cognition, it remains unclear what underlies individual variance on typical decision-making tasks.

**Methods:** We examined whether individual variance among Sprague-Dawley rats on a discounting decision-making task correlates with individual variance on memory and cognitive flexibility performance. Each task is conducted on an automated Figure-8 maze (cf., Graham et al., 2010). To assess decision-making ability, we used a probability-discounting task, in which animals choose between small/certain and large/uncertain rewards. The larger, uncertain reward had an expected value either 50% higher or 50% lower than the smaller, certain reward during “Uncertain Bias” and “Certain Bias” conditions, respectively. Over several days of exposure to each bias condition, rats biased their choices toward the reward with the higher expected value. We further evaluated decision-making performance based on two variables: risk preference (average preference of the large reward under high and low probabilities) and discrimination (the difference in preference of the large reward between high and low probabilities). We correlated individual performance across several measures to determine whether individual variance in performance on a probability-discounting task could be explained by variance in either working memory or cognitive flexibility.

**Results:** Probability discrimination did not correlate with individual risk preference. Accuracy on a delayed alternation working memory task negatively correlated with probability discrimination, but was unrelated to individual risk preference. Reversal learning performance positively correlated with risk preference, but not discrimination, on the probability-discounting task. However, a different measure of cognitive flexibility, the rate of response following a probability shift, did not correlate with either discrimination or risk preference.

**Conclusions:** The cognitive processes required for probability discrimination on a decision-making task may be related to those required for working memory performance, and risk preference may modulate reversal learning performance. Probability discrimination and risk preference both vary among individual rats and are distinct measures of decision-making cognition. It remains unclear whether they represent lower-level forms of cognitive processing that subsequently influence other higher-order tasks (e.g. working memory, reversal learning) or whether there may be additional factors that simultaneously influence each of these particular measures of performance.

**Acknowledgements:** This work was supported by NIH grant R01 MH64457 and the University of Washington Royalty Research Fellowship to JJK.

## **BOLD correlates of evidence integration during value-based decision making**

Ryan Jessup<sup>1\*</sup>, Antonio Rangel<sup>2</sup> and John P. O'Doherty<sup>1,2</sup>

<sup>1</sup>Institute of Neuroscience, Trinity College Dublin, Dublin 2, Ireland; <sup>2</sup>California Institute of Technology, Pasadena, CA, USA  
jessupr@tcd.ie

Objective: Influential theories of choice posit that decisions are generated following an integration of evidence about different decision options. The aim of the present research was to determine whether evidence accumulation occurs in the brain during value-based decision making, and to identify the brain regions exhibiting this response profile. Additionally, we were interested in comparing different classes of accumulatory models in order to see which class better reflects neural signals.

Methods: Here we used fMRI to measure BOLD signals while 20 hungry subjects made decisions between pairs of food items which could be consumed later. On each trial options were shown in a repeated serial fashion until an item was selected. After completing the choice section, participants then provided buying prices and familiarity ratings for each of the food items and upon completion of the experiment were given up to two foods to consume, depending on both random selection and their preferences. We applied a variety of computational models of decision making that posit different types of evidence accumulation to the fMRI data and tested for regions showing correlations with the accumulation process engendered by these different models during choice. Mean predicted signal per trial was obtained by simulating the models using the best fitting parameter values. This mean predicted signal was then regressed on the BOLD data, separately for each model.

Results: The results revealed that multiple regions correlated with the model predicted accumulation signals, including regions of the posterior parietal cortex, dorsolateral prefrontal cortex, and ventromedial prefrontal cortex. The different model classes revealed both overlapping and unique activation areas.

Conclusions: These results can help to disentangle the specific computations performed by different brain regions during decision making as well as provide evidence concerning which class of model can better explain neural activity during value-based decision making.

Funding: This work is supported by a grant from Science Foundation Ireland to JOD.

**Title:** Shockingly conservative: Threat of shock increases risk aversion

**Authors:**

Peter Sokol-Hessner\* (NYU/Caltech)

Jeffrey Hamilton (NYU)

Colin Camerer (Caltech)

Elizabeth Phelps (NYU)

\*Corresponding Author: hessner@gmail.com

**Abstract:**

Emotions have important roles in decision behavior as demonstrated by evidence both from studies that illustrate the influence of nominally irrelevant emotional manipulations upon choices, as well as those that show correlations of emotional variables with decision-making. In one example of the latter, a recent study showed that the relative physiological arousal responses to gain and loss outcomes correlated with individuals' estimated loss aversion (Sokol-Hessner et al, 2009). However, because arousal was not experimentally manipulated, any conclusions from these data are limited to statements of correlation. In order to move beyond such limitations, in this study participants faced a similar series of monetary choices in one of two contexts: "Threat", in which participants experienced intermittently reinforced auditory conditioning with a tone predicting electric shock, and "Safe", in which neither tone nor shock was present. Conditioning is a robust and well-characterized technique for eliciting discrete arousal responses, similar to those observed during the decision task in Sokol-Hessner et al (2009). Beyond the unique introduction of a shock-based conditioning paradigm, a three-parameter model of participants' choice behavior was also quantitatively estimated from individuals' choices, enabling improved specificity in isolating effects of the manipulation on behavior, as compared to previous studies. The threat of shock was found to systematically bias individuals' behavior such that they demonstrated increased risk aversion, relative to their choices made not in the presence of shock, and this effect could be distinguished from changes in loss aversion or individuals' consistency over choices. By combining conditioning with econometric models of decision behavior, these results contribute strongly to our understanding of how emotional arousal may causally affect the choices we make.

## **Chronic marijuana use is associated with distinct neural activity during monetary decision-making**

Jatin G. Vaidya<sup>1a</sup>, Robert I. Block<sup>2</sup>, Daniel S. O'Leary<sup>1</sup>, Laura B. Ponto<sup>3</sup>, Mohamed M. Ghoneim<sup>2</sup>, & Antoine Bechara<sup>4</sup>

<sup>1</sup>Department of Psychiatry, University of Iowa Hospitals & Clinics; <sup>2</sup>Department of Anesthesiology, University of Iowa Hospitals & Clinics; <sup>3</sup>PET Center, Department of Radiology, University of Iowa Hospitals & Clinics; <sup>4</sup>Brain and Creativity Institute, University of Southern California

<sup>a</sup>Correspondence at: [jatin-vaidya@uiowa.edu](mailto:jatin-vaidya@uiowa.edu)

**Objective.** Marijuana (MJ) acutely acts on cannabinoid receptors that are found in numerous brain regions including those involved in reward processing and decision-making. However, it remains unclear how long-term, chronic MJ use alters reward-based decision-making. In the present study, using [<sup>15</sup>O]water PET imaging, we measured brain activity in chronic marijuana users (CMJUs) and non-using control participants while they took part in the Iowa Gambling Task, a monetary decision making task that strongly relies on a neural circuitry in which the ventromedial prefrontal cortex is a key component for implementing advantageous decision-making.

**Methods.** Forty-six CMJUs and 34 control participants took part in the study. During PET imaging, participants took part in the standard version and a control version of the Iowa Gambling Task, as well as other tests. In the standard version of the task, participants must choose from one of four decks of cards. Participants learn through trial and error that two of the decks are advantageous (gains outweigh losses) and two are disadvantageous (losses outweigh gains). The task was timed so that the primary learning phase (trials 11 – 50) coincided with PET image acquisition. The control version of the task did not require decision-making; participants simply chose from each of the four decks in a pre-determined sequence.

**Results.** CMJUs and control subjects showed a great deal of common activity in regions involved in reward processing and decision-making, including the ventromedial prefrontal cortex and anterior cingulate cortex. In the two-group comparison, CMJUs did not show lower activity in the ventromedial prefrontal cortex as expected; in fact, using a more liberal threshold, CMJUs showed higher activity in this region. CMJUs also showed higher activity than control subjects in the cerebellum. Finally, greater duration of MJ use was associated with lower task-related activity in the anterior cingulate cortex.

**Conclusions.** CMJUs tend to rely more heavily on neural circuitry involved in decision-making and reward processing (ventromedial prefrontal cortex) and probabilistic learning (cerebellum) possibly reflecting inefficient information processing within these regions. Furthermore, given that duration of MJ use was associated with lower activity in the anterior cingulate, a region implicated in attention and error detection, chronic MJ use may alter the ability to modify decision making based on unexpected patterns of rewards and punishments.

### **Acknowledgments:**

This study was funded by the NIH, grant # 5R01DA019338.

## Effector-Specific Reward Value Updating in the Posterior Parietal Cortex

Y.T. Wong,<sup>1</sup> M.M. Fabiszak,<sup>1</sup> Nathaniel D. Daw,<sup>1,2</sup> and B. Pesaran.<sup>1\*</sup>

<sup>1</sup>Center for Neural Science, and <sup>2</sup>Department of Psychology, New York University.

\*Correspondence at: [bijan@nyu.edu](mailto:bijan@nyu.edu)

Objective: To make effective decisions we must first learn then update the values of the objects under consideration. How we learn values and then make choices for different effectors poses an interesting question. We tested the hypothesis that a single decision system guides choices for multiple effectors or alternatively whether separable systems guide choices for different effectors.

Methods: Two adult Rhesus macaques were trained to perform a variant of a two armed-bandit task. Animals chose between two targets each associated with a different reward value. Choices were made with either a single effector (reach-only, or saccade-only), or with two effectors (reach-and-saccade). The effector used on each trial was interleaved randomly, however the reward distributions of the two targets was kept constant across trials allowing the reward values to be learnt from either effector. The mean of the reward distributions were kept consistent for 50-60 trials after which were changed in an un signaled manner. Individual neurons were recorded from the posterior parietal cortex as the animals performed this task. Seventy neurons were recorded from the posterior reach region (PRR), an area active before reaches and 47 neurons from lateral intraparietal area (area LIP), an area active before saccades.

Results: Firing rates of neurons in both area LIP and PRR were modulated by reward values however area LIP showed stronger modulations than PRR. Choice selectivity emerged significantly faster in area LIP ( $579 \pm 41$  ms) than PRR ( $691 \pm 29$  ms, (mean  $\pm$  s.e.m)  $p=0.04$ ; Rank Sum test). When comparing the build up of choice information before and after the reward transition, PRR showed a significant slowing of the build up after the transition which we attribute to the learning of the new reward values. This difference was not seen in area LIP indicating differences in how reward values are updated in the two areas. To analyze how the animals learned reward values from rewards received from different effectors, a model that explained choices as a linear combination of previous rewards was fit. This model showed that a higher waiting was given to rewards received on trials in which a saccade was involved.

Conclusions: These results suggest that separate value updating rules may be applied to rewards received from different effector movements and that separable decision systems may guide movements for different effectors.

Acknowledgements: This study was funded by the NSF-NIH CRCNS MH087882, Alfred P Sloan Foundation, McKnight Foundation, Burroughs Wellcome Fund, and NYSTAR.

# A Neural Substrate of Probabilistic and Intertemporal Choice Within a Single fMRI Experiment

Hiroyasu Yoneda <sup>\*†‡</sup>      Sobei H. Oda <sup>§</sup>

**Objective:** The purpose of our study is to investigate differences and similarities in brain activities between choice under uncertainty and intertemporal choice. Some researchers have found that risk preference are correlated strongly with delay preference across individuals; others claim that the correlations are weak or absent. We did an fMRI experiment to examine risky choice and intertemporal choice within the same framework.

**Methods:** Twenty one subjects participated in our study. In an MRI they were repeatedly asked two-option questions (90 questions for each subject). Three conditions were used in the task: Delay tasks (choices between future and present/nearer future rewards) Risk tasks (choices between risky and certain/less risky rewards) and Control tasks (trivial delay or risk tasks).

**Results:** We found that precuneus and parahippocampal gyrus were activated more when they chose a future reward (1 or 2 weeks later) than they were when an uncertain alternative (40% or 80%) is chosen. In contrast, Orbitofrontal cortex (OFC) and angular gyrus were activated more when they chose a uncertain alternative than they were when an future reward is chosen.

**Conclusions:** Together with other observations, our results support, at least is consistent with, the self-projection theory proposed by Buckner and Carroll (2006): people imagine future themselves to choose a future reward. In addition, we examined our subjects' brain activities according to their preferences: certainty preference, risky preference, early-reward preference, and delay-reward preference, which give other evidences to support the theory.

**Acknowledgements:** This study was supported by Japan Society for the Promotion of Science (JSPS) and Open Research Centre by the Ministry of Education, Culture, Sport, Science and Technology (FY2006—FY2008). And this research owes much to the thoughtful and helpful comments of Jiro Okuda (Kyoto Sangyo University).

---

\*Research Fellow of the Japan Society for the Promotion of Science

†Institute of Social and Economic Research (ISER), Osaka University

‡Correspondence at: [hiroyasu.yoneda@gmail.com](mailto:hiroyasu.yoneda@gmail.com)

§Kyoto Sangyo Univeristy

## Working Memory and Intertemporal Choice

Sarah J. Getz<sup>1,2</sup>, Damon Tomlin<sup>1,2</sup>, Leigh E. Nystrom<sup>1,2</sup>, Jonathan D. Cohen<sup>1,2</sup> and Andrew R. A. Conway<sup>1,2</sup>

<sup>1</sup>Princeton University Department of Psychology; <sup>2</sup>Princeton Neuroscience Institute

Correspondence at: [sgetz@princeton.edu](mailto:sgetz@princeton.edu)

**Abstract:** Intertemporal choice (ITC) involves trade-offs between costs and benefits at different points in time. Recent neuroimaging data suggest that ITC involves a dynamic interplay between brain regions associated with cognitive control and working memory (WM), such as prefrontal cortex, and structures in the limbic system associated with the evaluation of reward. Here we test the hypothesis that engaging participants in an unrelated task that competes for the cognitive control mechanisms of WM during ITC should lead to more impulsive decision-making. Previous research on the effect of WM load on ITC is inconclusive on this point because it has been difficult to distinguish impulsive from random responding (Franco-Watkins, Pashler, & Rickard, 2003; Hinson, Jameson, & Whitney, 2003). Participants in this experiment made ITC decisions while simultaneously performing an n-back task under high load (3-back) or low load (0-back). A stair-casing method was used to determine each individual participant's indifference point and subsequent choice values were titrated around this point. Participants in the high load condition chose the sooner reward more often than those in the low load condition and participants in both groups demonstrated a reliable pattern of responding around the indifference point. These results support the conclusion that WM working memory load results in more impulsive decision-making.

## Neural Correlates of Cognitive Dissonance and Choice-Induced Preference Change

Keise Izuma<sup>1,2\*</sup>, Madoka Matsumoto<sup>1</sup>, Kou Murayama<sup>3</sup>, Kazuyuki Samejima<sup>1</sup>, Norihiro Sadato<sup>4</sup>, Kenji Matsumoto<sup>1</sup>

<sup>1</sup>Brain Science Institute, Tamagawa University, Tokyo, Japan. <sup>2</sup>Current address; Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, USA.

<sup>3</sup>Department of Psychology, University of Munich, Munich, Germany. <sup>4</sup>Department of Cerebral Research, National Institute for Physiological Sciences (NIPS), Aichi, Japan.

\*Correspondence at: [izuma@caltech.edu](mailto:izuma@caltech.edu)

**Objective:** While economic theory views that actions are simply reflective of individuals' preferences, a psychological theory claims that actions can create preference. According to the cognitive dissonance theory, after making a difficult choice between two equally preferred items, the act of rejecting a favorite item induces uncomfortable feeling called "cognitive dissonance," which in turn motivates individuals to change their preferences to match their past decisions (i.e., reducing preference for rejected items). However, neural processes underlying dissonance-induced preference change remain largely unknown. We conducted an fMRI experiment to examine how subjects' preferences for the same foods changed after making difficult choices between two preferred foods and also how the brain responds to cognitive dissonance (i.e., discrepancy between their past decisions and preferences).

**Methods:** Images of one hundred sixty food items (e.g., chips, chocolate, etc) were used in the experiment. The experiments consisted of four parts; 1) Preference task 1, 2) Choice task, 3) Preference task 2, and 4) Post-Experimental Choice task. Except for the Post-Experimental Choice task, all tasks were performed during fMRI scanning.

**Results:** We first found that subjects' reported preferences during the Preference task 1 were positively correlated with the activity in the anterior striatum. Then, as predicted, the striatal activity was significantly reduced from the Preference task 1 to the Preference task 2 for the preferred foods which were rejected during the Choice task compared to other control conditions. Furthermore, we found that the dorsal anterior cingulate cortex (dACC) and dorsolateral prefrontal cortex (DLPFC) tracked the degree of cognitive dissonance on a trial-by-trial basis.

**Conclusions:** This study provides evidence that the mere act of making choices can alter individual's self-report preference as well as its neural representation, and dissonance-induced preference change recruits the same neural network underlying the monitoring of lower level conflict and the subsequent implementation of control.

### Acknowledgements:

Supported by a Grant-in-Aid for JSPS Fellows to K. Izuma, a Grant-in-Aid for Scientific Research C#21530773 to K. Matsumoto and Tamagawa University Global Center of Excellence grant from the Japanese government.

Rats respond to the opponents' change in strategy in a competitive game

Tassi, LE and Xavier, GF

Department of Physiology, Biosciences Institute, University of S Paulo, S Paulo, Brazil

Correspondence at: [letassi@terra.com.br](mailto:letassi@terra.com.br)

**Objective:** We studied whether Wistar rats were capable of playing the Matching Pennies game (MPG), and how they responded to the use of different strategies by a computer opponent. MPG is a simple strategic game, in which players' choices, to be optimal, must be independent of previous choices and payoffs. According to reinforcement learning models, adjustments of choice policy are based on discrepancies between expected and obtained rewards, thus generating more predictable choice sequences, which lead to a lower reward rate.

**Methods:** Eleven male Wistar rats were trained to play a nose-poke version of the game. In each trial, after nose poking a central hole, they had to choose one of two lateral holes. When the animal and the computer chose the same hole, reward (sucrose solution) was provided. The computer was programmed to use two different algorithms to predict the animals' next choice, and thus exploit eventual biases in their choice sequence. The first algorithm ONLY exploited statistical biases present in the right (R)-left (L) choice sequence, e.g. the probability of L choice after a RRL sequence. So, a reinforcement learning policy of repeating rewarded choices was not penalized. In the second algorithm, the sequence of choices and the payoffs was taken into account, e.g. the probability of a R choice after a R(+)R(-)L(+) sequence (where (+) and (-) stand for rewarded and unrewarded choices, respectively). Consequently, in the second algorithm, the probability of getting rewarded was smaller when subjects used the "win-stay-lose-shift" strategy (WSLS). Therefore, if rats were sensitive to the opponents' strategy, their probability of using a WSLS in the second algorithm would be smaller.

**Results:** When playing against algorithm 2, subjects used WSLS significantly less than when playing against algorithm 1 (Wilcoxon signed-rank test,  $p=0.002$ ). There was no significant difference in the reward rate (Wilcoxon signed-rank test,  $p > 0.05$ ), which was 48%, significantly different from the expected optimal rate of 50% (Wilcoxon signed-rank test,  $p = 0.0001$ ).

**Discussion:** Rats were capable of approaching the optimal strategy in the MPG, and were sensitive to changes in the opponent's strategy, responding accordingly. The changes in the probability of WSLS and the lower reward rate, demonstrate that reinforcement learning is at the core of the animals' strategy.

**Support:** FAPESP, CAPES e CNPq

# Exposure to Economic Arguments Reduces Delay Discounting

Nicole Senecal\* and Joseph W. Kable  
Department of Psychology, University of Pennsylvania  
\*Correspondence at: senecaln@mail.med.upenn.edu

## Objective

When making decisions that involve tradeoffs between the quality and timing of desirable outcomes, people consistently discount the value of future outcomes. Discounting is often measured in the laboratory using monetary rewards, by asking subjects to choose between a smaller amount of money available immediately and a larger amount available after a delay. The high degree of discounting observed on this task is extremely puzzling to many economists. For decisions about money, a normative argument can be made that regardless of personal time preferences, subjects should only turn down rates of return from the experimenter that are lower than rates available to them elsewhere, such as on the market. However, all studies of delay discounting cite monetary discount rates that are orders of magnitude higher than market interest rates. Here we ask whether one reason that subjects discount to such a degree is that they are simply unaware of the normative argument or how it applies to the task.

## Methods

Subjects read a “financial education guide,” which explicitly outlined the normative argument with concrete examples. Each subject completed a monetary delay discounting task with real financial incentives at three time points: before, immediately after, and one month after reading the guide. When returning one month later, subjects’ sessions were run by a different experimenter, with no mention made of the material read at the previous session.

## Results

Immediately after reading the guide, discount rates decreased significantly. One month later, discount rates had increased, but remained significantly lower than before subjects had read the guide. This decrease was not due to subjects adopting simple decision rules, such as setting a cutoff amount or always selecting the later reward. Rather, subjects appeared to integrate the information from the guide and adjust their degree of discounting.

## Conclusions

These results demonstrate that one reason for high monetary discount rates is that subjects have not considered relevant economic arguments, although other factors (i.e., liquidity constraints, uncertainty regarding payment) are also needed to fully explain the deviation from market interest rates. These results also show how a simple educational intervention can have a significant and long-lasting effect on decision-making. As such, this paradigm should prove useful for studying the neural mechanisms by which education can change behavior. Ongoing experiments are examining this question.

## **Subliminal Brand Priming Influences Incidental Decision-Making**

Philip G. Harris<sup>1</sup>, Carsten Murawski<sup>2\*</sup>, Stefan Bode<sup>3</sup>, Juan F. Domínguez D.<sup>4</sup> and Gary F. Egan<sup>4,5</sup>

<sup>1</sup> *Department of Management and Marketing, The University of Melbourne, Melbourne, Australia*

<sup>2</sup> *Department of Finance, The University of Melbourne, Melbourne, Australia*

<sup>3</sup> *Max-Planck-Institute for Human Cognitive and Brain Sciences, Leipzig, Germany*

<sup>4</sup> *Howard Florey Institute, Melbourne, Australia*

<sup>5</sup> *Centre for Neuroscience, The University of Melbourne, Australia*

\*Correspondence at: carstenm@unimelb.edu.au

**Objective:** In the present study we examined incidental effects of brand presentation on financial decision-making. We explored the possibility that brand images with affective value impact on decisions by influencing affective systems underlying decision-making.

**Methods:** Sixteen participants attended a behavioural session and a subsequent fMRI session. Participants performed a temporal discounting task (see Kable & Glimcher, 2005), and a perceptual direction task (control). In the temporal discounting task, half of the choices were subliminally primed (16ms, pre- and post-masked) with an affective image (“Apple” logo), half with a neutral image (cup). We then used multivariate pattern classification to predict different aspects of decisions for each stimulation group separately from local spatial brain activation patterns, using an unbiased, moving “searchlight” (Kriegeskorte et al., 2006; Haynes et al., 2007). Additionally, we used univariate analyses to search for brain regions in which activation was directly modulated by subjective values of choices in each trial.

**Results:** The behavioral experiment demonstrated that priming with the subliminally presented “Apple” logo shifted participants’ choices relative to the neutral condition. Multivariate pattern classification showed that activation patterns in orbitofrontal cortex (OFC) predicted the general task type (temporal discounting vs. control). We could further predict which subliminal prime was used within the temporal discounting condition from activation patterns in anterior cingulate cortex (ACC) / medial prefrontal cortex (PFC). Decision outcomes (now vs. later) were also differentially encoded for each priming group, with predictive patterns in OFC / medial PFC for “Apple” and inferior parietal lobe for “neutral” (cup). Univariate analyses showed that for the temporal discounting task, activation in striatum, dorsolateral PFC and ACC was correlated with subjective value of the choice.

**Conclusions:** We could show a subliminal priming effect of an emotionally-relevant stimulus on an unrelated temporal discounting task. The stimulus seemed to modulate valuation of the incidental reward, as indicated by differential activation patterns in brain regions previously implicated in valuation (striatum, medial PFC). Overall, subjective values were directly reflected in activation in brain areas, which are part of a network underlying temporal discounting and, more broadly, affective learning and decision making processes (Kable and Glimcher 2007; McClure et al. 2004; Singer et al. 2009).

## Frontopolar cortex contributes to choice exploration by tracking recent payoff trends

Christopher K. Kovach<sup>1,2</sup>, Nathaniel Daw<sup>3</sup>, David Rudrauf<sup>2</sup>, Daniel Tranel<sup>2</sup>, John P. O'Doherty<sup>1,4</sup>,  
Ralph Adolphs<sup>1</sup>

<sup>1</sup>Division of Humanities and Social Sciences, California Institute of Technology

<sup>2</sup>Department of Neurology, The University of Iowa

<sup>3</sup>Center for Neural Science and Department of Psychology, New York University

<sup>4</sup>Trinity College Institute of Neuroscience, Trinity College, Dublin, Ireland

**Objective:** The functions of frontopolar cortex (FPC) have long been enigmatic. Recent neuroimaging studies find that it is activated during exploring and switching in the tracking of variable rewards during n-armed bandit tasks, but leave unresolved its causal contributions to behavior. We examined the ability of patients with brain lesions that include frontopolar cortex to track variable payoffs in a 4-armed bandit task previously shown to elicit BOLD activity in FPC.

**Methods:** 6 patients with lesions encompassing FPC (FP group), 8 brain-damaged comparison subjects whose lesions spared FPC (BDC group), and 14 normal comparison subjects (NC group) participated in a 4-armed bandit task for a fictive monetary reward. Mean payoff generated by each bandit followed a random walk. Subject choices were fitted to two different models: (1) a multinomial logistic (MNL) model in which the predictor was the history of the previous 5 payoffs for each bandit, (2) a reinforcement learning model augmented with the difference between the last two payoffs (RL+LD). Responses from each subject were fitted to both models. Maximum likelihood parameter estimates were compared for each group using non-parametric statistics.

**Results:** Significant group differences were observed in the lag-1 (Kruskall-Wallis  $P = 0.035$ ) and lag-2 ( $P = 0.065$ ) regression weights under model 1. Of note, NC subjects showed a negative weighting on the lag 2 payoff (Wilcoxon rank-sum  $P = 0.025$ ), which contradicts the geometric weighting predicted by an RL model. We hypothesized that the negative lag-2 weight reflected a tendency to extrapolate trends. We therefore fitted an RL model augmented with a linear dependence on the difference between the last two payoffs (RL+LD). Groups showed no significant difference in the parameters of the RL component (learning rate:  $P = 0.45$ , softmax temperature:  $P = 0.65$ ), whereas the group effect for the LD term was significant ( $P = 0.0067$ ). Post-hoc tests revealed that BDC and NC groups significantly weighted the lag-difference ( $P = 0.0017$  and  $P = 0.016$ , respectively), whereas FP subjects did not ( $P = 0.56$ ). FP subjects differed from NC and BDC (Tukey's HSD,  $P = 0.051$  and  $P = 0.0050$ , respectively), while NC and BDC groups did not differ ( $P = 0.43$ ).

**Conclusion:** Using model-based analyses of choice behavior in patients with focal brain lesions, we show that the frontopolar cortex makes a specific contribution to exploratory behavior: extrapolating current trends from comparisons among the most recent outcomes. Moreover, the contributions of trend following and reinforcement learning to choice valuation are functionally dissociable.