Vasopressin increases human risky cooperation by enhancing the functional coupling of the dlPFC and pallidum

C. Brunnlieb [1,2,3,4*], G. Nave [5], C.F. Camerer [5,6], B. Vogt [3,4], T.F. Münte [1], M. Heldmann [1,2]

[1] Department of Neurology, University Lübeck; [2] Center for Behavioral Brain Sciences, Magdeburg; [3] Department of Economics, Otto-von-Guericke University Magdeburg; [4] Institute for Social Medicine, Epidemiology and Health Economics, Otto-von-Guericke University Magdeburg; [5] Department of Computation and Neural Systems, California Institute of Technology; [6] Department of Humanities and Social Sciences, California Institute of Technology

*Correspondence to: claudia.brunnlieb@ovgu.de

Objective: The neuropeptide arginine vasopressin (AVP) is known to mediate complex social behavior like pair bonding, social recognition, and aggressive behavior in mammals. In order to investigate AVP's impact on human behavior and related neural structures in risky social cooperation, we recorded functional magnetic resonance imaging (fMRI) while participants performed stag hunt (SH) games with varying cooperation incentive levels. We hypothesized that AVP would only increase cooperative behavior when the incentive to cooperate is high.

Methods: In a double-blind placebo-controlled fMRI-study, 34 healthy male participants received either AVP (20 IU) or a Placebo intranasally and played variations of a 2x2 SH game with a fixed partner (no feedback). In the game, players chose between a cooperative (risky) strategy and a non-cooperative (safer) strategy. 105 SH games were derived from 7 basis-games that differed only in the incentive to cooperate.

Results: AVP treatment increased cooperative behavior when the incentive to cooperate was high. The AVP-related increase in cooperation was mediated by two neural mechanisms: During cooperative (risky) choices (1) AVP down-regulated the activity in the left dorsolateral prefrontal cortex (dlPFC), a brain region known to be activated during tasks that require cognitive control and increased mental effort; (2) AVP strengthened the functional coupling between the left dlPFC and the left pallidum, a region of the reward circuitry with many AVP receptors.

Conclusions: Our findings suggest that under AVP treatment, cooperation requires less cognitive control and emotional regulation possibly due to a decrease in the perception of social risk. Furthermore, it appears that AVP increases risky cooperation by alleviating aversion to social risk which might be indicated by the observed dlPFC coupling with reward signals in the pallidum.

Acknowledgements: This work was supported by a special grant of the Center for Behavioral Brain Sciences, Magdeburg (MH, CB), Deutsche Forschungsgesellschaft (TM), and the Betty and Gordon Moore Foundation (CC, GN).

Functional brain connectivity identifies human motives

Grit Hein[1], Yosuke Morishima[2], Susanne Leiberg[1], Sunhae Sul[3], and Ernst Fehr[1]

[1] Laboratory for Social and Neural Systems Research, Department of Economics, University of Zurich, Switzerland
[2] Department of Psychiatric Neurophysiology, University Hospital of Psychiatry, University of Bern, Switzerland
[3] Department of Psychology, Korea University, South Korea

[*]Correspondence at: ernst.fehr@econ.uzh.ch

Objective: Almost any given human behavior can be driven by different underlying motives. Thus, by merely observing behavior it may be impossible to identify the motive that drives behavior. Here we examine whether data on functional brain activity or connectivity can be used to identify human motives in situations where purely behavioral observation necessarily fails to achieve this because the different motives have behaviorally identical implications. As a byproduct, our analysis also identifies the neural architecture of important human motives behind prosocial behavior.

Methods: In the fMRI scanner, participants made the same set of prosocial decisions either driven by an empathy or a reciprocity motive that we controlled experimentally. First, we identified neural regions involved in prosocial decisions. Second, we used Dynamic Causal Modelling to determine the individuals' architecture of neural interactions between these regions. Third, we submitted the individual DCM parameters to a classification algorithm (Support Vector Machine, SVM).

Results: The functional connectivity patterns – but not the activity patterns – captured by individuals' DCM parameters predict the two underlying motives with high accuracy, thus identifying whether a person's decision is driven by empathy or reciprocity. Our results further indicate that the neural circuitry that underlies the empathy motive is basically identical to the neural circuitry of basic prosocial motivations among anonymous strangers, suggesting that basic prosociality is likely to be driven by empathy. This finding contributes to the solution of a long-standing unresolved question in altruism research.

Conclusions: The distinct neural architecture of human motives enables the identification of motivational drivers of behavior when behavioral methods necessarily fail to do so. We show this in the domain of human prosocial decisions and – as an important byproduct – we identify the neural architecture behind two key drivers of human prosociality – empathy and reciprocity.

How serotonin and dopamine shape moral decision-making

Crockett MJ*[1,2], Siegel JZ[2], Kurth-Nelson Z[2], Ousdal OT[2], Story GW[2], Dayan P[3], Dolan RJ[2]

[1]Department of Experimental Psychology, University of Oxford
[2]Wellcome Trust Centre for Neuroimaging, University College London
[3]Gatsby Computational Neuroscience Unit, University College London

*Correspondence at: mollycrockett@gmail.com

Objective: How we evaluate the suffering of others is a central concern in moral decision-making, but its neural basis is unknown. Here, we investigated how people value others' pain relative to their own pain, and how serotonin and dopamine influence aversion to pain for self and others.

Methods: Two subjects participated in each experimental session under conditions of complete anonymity and were randomly assigned to the roles of 'Decider' and 'Receiver'. Deciders made a series of choices between a smaller amount of money plus a smaller number of mildly painful electric shocks, versus a larger amount of money plus a larger number of shocks. The Decider always received the money, but the shocks were allocated to the Decider on half of the trials and to the Receiver on the other half.

We deployed this paradigm in two behavioral studies (N = 39 and N = 41) and used a computational model of Deciders' choices to derive a pair of subject-specific *harm aversion* parameters that characterized the subjective cost of pain for self and others, respectively. We then carried out two double-blind, placebo-controlled pharmacological studies to investigate the effects of the serotonin reuptake inhibitor citalopram (N = 89) and the dopamine precursor levodopa (N = 86) on harm aversion for self and others.

Results: Across all four studies, we find that harm aversion for others is greater than harm aversion for self. In other words, most people will selflessly sacrifice more money to prevent others' pain than their own pain. Citalopram and levodopa had distinct and opposing effects on moral decision-making: citalopram increased harm aversion for both others and self, while levodopa selectively reduced harm aversion for others without affecting harm aversion for self. Crucially, neither drug influenced the physical *perception* of pain, suggesting a direct influence of serotonin and dopamine on preferences.

Conclusions: We show that serotonin and dopamine exert distinct effects on moral decision-making by differentially modulating the valuation of pain for self and others. Our findings have implications for understanding antisocial behavior in psychiatric disorders associated with abnormal serotonergic and dopaminergic function.

Acknowledgements:
This study was funded by the Wellcome Trust (092217/Z/10/Z).

Eye movements reveal the effect of branding on consumer decisions

I. Krajbich[1,2], R. Ratcliff[1], C. Retzler[3], A. Rangel[4], M. Philiastides[3]

[1] Department of Psychology, The Ohio State University
[2] Department of Economics, The Ohio State University
[3]Institute of Neuroscience and Psychology, University of Glasgow
[4]Division of Humanities and Social Sciences, Caltech

Objective: Product branding is a crucial dimension of consumer choices. Recent work has suggested that branding information and subjective product preference may be integrated into a single source of evidence in the decision-making process. Here we investigate how exactly these two sources of information are combined, by employing the attentional drift-diffusion model (aDDM) to relate choices and reaction times to the relative gaze time on the two products and their brands.

Methods: We carried out an experiment in which subjects made a series of hypothetical preference decisions between two items of clothing paired with different designer brands. In control trials subjects also made preference-based clothing decisions, but with phase-scrambled brand images. While subjects made these choices, we tracked their eye-movements. Beforehand we also collected separate individual ratings for each clothing item and brand. We then used subjects' ratings and gaze patterns as inputs to the aDDM to test whether these measures alone could account for subjects' choices and reaction times.

Results: Using the aDDM we were able to accurately predict the influence of gaze time on the probability of choosing the left or right item. Comparing the intact brand trials to the scrambled control trials, we find that subjects spent more time looking at the brand information, took longer to make their decisions, and were more likely to choose an item if it was paired with a preferred brand. Furthermore, we were able to use the aggregate fraction of time spent looking at the brands to predict the average influence of the brand ratings on subjects' choices. This relationship was further established with a significant across-subject correlation between brand gaze time and brand weight in their utility functions. Finally, consistent with previous aDDM findings, we observed no correlation between item or brand ratings and gaze duration.

Conclusions: Our results indicate that branding information and subjective product preference are combined together in a multi-attribute drift-diffusion model, where the relative weights on the two attributes are determined by the gaze time on the product vs. brand. These findings lend further support to the aDDM as a common mechanism underlying value-based decisions and are consistent with the hypothesis that in binary choice, attention leads to preference, and not vice-versa.

The Origin of the Pain of Paying: Evidence from fMRI and behavioral experiments

Nina Mazar[1], Hilke Plassmann[2,3*], Nicole Robitaille[1], Axel Linder[4]

[1]University of Toronto, Canada; [2]INSEAD, France; [3]Cognitive Neuroscience Unit INSERM & Ecole Normale Superieure, France; [4]Univerisity of Tübingen, Germany

[*]Correspondence at: hilke.plassmann@insead.edu

Objective: Behavioral decision theories suggest when purchasing a product, people experience hedonic competition between the anticipated pleasure derived from consuming the product and the anticipated displeasure derived from paying for it. The latter has been referred to as "pain of paying". Interestingly, however, our understanding of the "pain of paying" is still inconclusive. The goal of this paper is to shed further light on its psychological basis.

Methods & Results: In study 1 (N=19), we used functional magnetic resonance imaging to examine how the human brain processes different types of payment conditions, i.e. paying with money vs. paying by tolerating electric shocks. Although we found that purchase decisions involving money and shocks were similar on a behavioral level, we found significant differences at a neural level: Whereas paying with "pain" recruits brain areas involved in different aspects of pain processing including somatosensory and higher order cognitive aspects, paying with money affects brain areas involved in valuation and also in higher order cognitive aspects of pain processing.
In two behavioral follow-up studies, we then examined whether the pain of paying is experienced as a physical pain, psychological (higher order) pain, or whether it is not experienced as a pain at all. In study 2 (N = 109) we examined whether making pain more salient would influence consumer's WTP for products, and if so, which types of pain would have an influence. We found that participants primed with psychological pain-related words were willing to pay significantly less as compared to participants primed with neutral or physical pain-related words.
In study 3 (N = 173), participants were first given a placebo pill disguised as either a drug that decreased sensitivity to experienced pain (either physical or psychological), increased sensitivity to experienced pain (either physical or psychological), or a dietary supplement (control), after which we measured participants' WTP for an Amazon gift certificate. We found that indeed, participants' maximum WTP differed by the placebo condition they were in: Participants who were administered the psychological pain enhancer had a significantly higher maximum WTP than participants who were administered the psychological pain reliever. Participants who were administered the physical pain enhancer versus reliever did not differ from each other and the control, and their maximum WTP was between the two psychological pain conditions.

Conclusions: Taken together, across three studies we provided evidence for the idea that the pain of paying is a painful experience, albeit a psychological one.

Testosterone and Trading: Biological Driver of Mispricing

A. Nadler, [1*] P. Jiao, [2] V. Alexander, [2] C. Johnson, [3] and P. Zak[2,3]

[1]Richard Ivey School of Business at Western University; [2]Claremont Graduate University;
[3]Loma Linda University

[*]Correspondence at: anadler@ivey.ca

Objective: Hormones are chemical messengers used by the body to change the likelihood of behavior. The influence of hormones on financial decisions is largely unknown beyond correlational studies. Using treatment and placebo groups in experimental financial markets, we tested the hypothesis that testosterone drives competitive and "noisy" trading behavior that causes sustained mispricing ("bubbles") and higher price volatility among male traders.

Methods: One hundred and forty four males (called "traders") participated in groups in a study involving experimental financial markets. Trading groups were given either a testosterone or placebo gel prior to trading to test the influence of the hormone on their behavior and its attendant impact on the market. Traders earned money based on their trading decisions in this real-time market with real money earnings. The asset they traded had a fundamental value that was known to all traders throughout the experiment, allowing for clear identification of price "bubbles". Individual trading data was used to assess drivers of differential pricing and volatility levels between high- and average-testosterone trading sessions.

Results: Blood draws showed testosterone levels increased 63% on average in the treatment groups (which is within the range of normal variation) while placebo traders showed no statistical difference. We found that testosterone caused more competitive bidding that persisted longer, thereby driving differences in prices between testosterone compared to placebo trading groups. Price bubbles formed and grew in the testosterone sessions quickly, and dramatically crashed towards the end of trading rounds. Also, high-testosterone traders did not incorporate fundamental value in their trading decisions, while traders given placebo show strong evidence of adapting to the changing fundamental value. Survey results show high-testosterone traders considered themselves more "talented" and less "lucky" relative to placebo-treated traders. Money earned from trading negatively correlated with testosterone levels across treatment groups.

Conclusions: These results suggest testosterone has significant activational properties that affect men's financial trading with meaningful implications to market prices and volatility. These results show testosterone is a biological driver of mispricing in financial trading by the channel of competitive bidding and attenuated attention to intrinsic value of the underlying asset.

Acknowledgements:
This study was funded by IFREE and the Center for Neuroeconomics Studies

Stock Ownership and Learning from Financial Information

Sarah Rudorf[1], Camelia M. Kuhnen[2*] and Bernd Weber[1]

[1]University of Bonn, Germany; [2]University of North Carolina, Kenan-Flagler Business School, USA

[*]Correspondence at: camelia_kuhnen@kenan-flagler.unc.edu

Objective: An implicit assumption in economics is that market participants are able to learn the same way from new information, irrespective of the composition of their portfolio. While theoretical work has shown that previous portfolio choices may influence investors' *utility function (1)*, it is possible that these prior choices might also change investors' *beliefs* or the learning rules they use to incorporate financial market news. Here, we study whether stock ownership status changes people's ability to use financial information correctly and investigate the brain mechanisms underlying this effect.

Methods: We collected behavioral and fMRI data in a sample of adults (*N*=46, ages 29-49 yrs) performing an investment choice task *(2)* where they had to select during 96 trials one of two assets: a risky stock about which there was uncertainty whether or not it paid dividends from a good or a bad distribution, and a safe bond with known payoffs.

Results: Prior investment decisions interfere with people's ability to correctly update their beliefs about the distribution of the stock's payoffs. If people's most recent choice is the stock, they will update beliefs more after observing a high dividend, rather than a low dividend. If their most recent choice is the bond, people will update their beliefs more after observing a low dividend of the stock, rather than a high one. Thus, erroneously, investors learn more from new information which ex-post justifies their prior portfolio choice.

This behavioral effect is driven by an asymmetry caused by prior choices in the response to new information observed in the anterior cingulate cortex (ACC), the ventromedial prefrontal cortex (VMPFC) and the striatum. When people choose the stock, activation in the ACC, VMPFC and striatum is higher if the new dividend paid by the stock is high, rather than low. When people choose the bond, there is less sensitivity in activation in these areas to high versus low stock dividends. Bond holders' muted brain reactivity to new stock information predicts the errors in beliefs subsequently expressed by these individuals regarding the stock payoff distribution and their willingness to invest in the stock in future trials.

Conclusions: Our findings can help explain two financial markets puzzles. First, the fact that stock holders do not update sufficiently from low dividends can explain the disposition effect *(3)*, i.e., that investors are reluctant to sell stocks that have not performed well. Second, that fact that bondholders do not update sufficiently from high dividends, and thus are overly pessimistic about future stock payoffs, can help explain why most households do not participate in the stock market *(4)*.

1. N. Barberis, M. Huang, T. Santos, *The Quarterly Journal of Economics* (2001).
2. C. M. Kuhnen, *Journal of Finance* (forthcoming).
3. T. Odean, *The Journal of Finance* **53** (1998).
4. J. Campbell, *Journal of Finance* **61**, 1553 (2006).

Imaging brain-wide responses to stimulation of VTA dopaminergic neurons with ofMRI

E. Ferenczi,[1,2] K. Zalocusky, [1,2] K. Katovich,[3] L. Grosenick,[1,2] G. Glover, [4]
K. Deisseroth,[1,2,5] and B. Knutson[2,3*]

[1]Bioengineering; [2]Neurosciences; [3]Psychology; [4]Radiology; [5]Psychiatry;
Stanford University, Stanford CA

[*]Correspondence: knutson@psych.stanford.edu

Objective: We describe the combined use of optogenetics and functional magnetic resonance imaging (ofMRI) to map whole-brain spatiotemporal patterns of Blood Oxygen Level Dependent (BOLD) activity during stimulation of midbrain dopaminergic neurons in awake rodents. Dopamine neurons play a central role in reward processing and when stimulated directly have been shown to support self-stimulation behavior in rodents (Witten et al, 2010). Human neuroimaging studies have shown that reward cues can increase the BOLD signal in mesolimbic regions, including the nucleus accumbens (O'Doherty 2004, Knutson & Cooper 2005) however the connection between this BOLD response and dopaminergic neuron activity has not yet been causally demonstrated.

Methods: We assessed the contribution of midbrain dopaminergic neurons to the brain-wide BOLD response using a causal manipulation that is known to drive reward-seeking behavior. We expressed cre-dependent channelrhodopsin-2 (ChR2) in the midbrain of transgenic (TH-cre) rats, targeting dopamine neurons in the ventral tegmental area (VTA) and substantia nigra (SN). ChR2-expressing rats quickly learned to press a lever to deliver 470 nm light to the midbrain, compared to YFP control rats and to an inactive lever. Awake functional MRI scanning (7 Tesla magnet) was enabled through a 5-10 day habituation protocol to accustom the rats to the MRI environment, in combination with MRI-compatible head fixation and pulse sequence optimization for rapid image acquisition.

Results: Midbrain stimulation (2 s 470 nm light, 20 Hz, 6 mW at fiber tip) generated a robust BOLD response in the ipsilateral dorsal and ventral striatum ($n = 6$ rats, p-value < 0.01 corrected for multiple comparisons). The ventral striatal response was positively correlated with reward-seeking behavior. Pharmacological experiments were performed on a subset of subjects and demonstrated that the striatal BOLD response was susceptible to blockade by dopamine receptor antagonists, which was reversible on drug washout.

Conclusions: These findings have implications for the interpretation of reward-related fMRI tasks in humans, for determining the impact of pharmacological agents on dopamine signaling across brain regions, and for understanding pathological conditions in which dopamine signaling and reward processing are disrupted, such as depression and addiction.

Acknowledgements:
EF is supported by Fulbright, IIE and SGF, KZ is supported by the NSF GRFP, KK and BK are supported by BioX, GG is supported by NIH and KD is supported by HHMI, NIH, NSF, DARPA, and SFARI.

State prediction errors in the orbitofrontal cortex

S.C.Y. Chan,[1*] N.W. Schuck[1], N. Lopatina, [2] and Y. Niv. [1]

[1]Princeton Neuroscience Institute, Princeton University;
[2]Program in Neuroscience, University of Maryland.

[*]Correspondence at: scychan@princeton.edu

Objective: How do we learn the transition structure of the world? One plausible algorithm is via "state prediction errors" that encapsulate the discrepancy between the state that we expect and what we actually observe. These state prediction errors (SPE) can be used to update our internal estimates of transition probabilities between states. Previous experiments indicated that the orbitofrontal cortex (OFC) is involved in the learning of transition structure, leading us to test the hypothesis that the OFC represents SPEs.

Methods: We scanned twenty four subjects (18-34 years old) using fMRI. Subjects observed stimulus-outcome pairs that were designed to elicit state prediction errors in the absence of value prediction errors. This was achieved by using different outcomes that had similar value (namely, M&M candies of different colors). We analyzed OFC activity using multi-voxel pattern analysis, in search of representation of the three necessary components of SPE: start state, end state, and magnitude of surprise.

Results: As predicted, we found that BOLD activity in the OFC can be used to decode start state, end state, and magnitude of surprise—the three components of SPE. Importantly, we did not find representation of SPE components in areas that had previously been implicated in representing SPEs (Glascher et al., 2010). These areas were found to encode only the magnitude of surprise, and thus to be more consistent with an attention signal rather than SPE.

Conclusions: Our results are consistent with the idea, suggested by previous experiments, that the OFC represents (and possibly computes) state prediction errors, which can be used to learn transition probabilities between states of the world.

Acknowledgements:
This study was funded by the NIH (grant # R01MH098861 and # T32MH065214), and the NSF/NIH Collaborative Research in Computational Neuroscience Program (grant # NSF IIS-1009542).

A dynamical model of normalization: the interaction between time, value, and choice

K. Louie[1*], T. LoFaro[2], R. Webb[1], and P.W. Glimcher[1,2]

[1]Center for Neural Science, New York University
[2]Department of Mathematics and Computer Science, Gustavus Adolphus College

[*]Correspondence at: klouie@cns.nyu.edu

Objective: During decision-making, divisive normalization implements a context-dependent value code in frontal and parietal cortices that explains significant behavioral violations of rational choice theory. However, while decision-making is a dynamic process with complex temporal properties, most models of normalization are time-independent and little is known about the time-varying interaction between normalization and choice. Here, we show that a simple dynamical system model of normalization produces characteristic value coding dynamics. These results suggest a specific circuit mechanism for value representation and predict novel speed-accuracy interactions in economic decision-making.

Methods: We constructed a simple differential equation-based dynamical firing rate model of a decision circuit characterized by recurrent inhibition. In this model, choice options are represented by paired excitatory (*R*) and inhibitory (*G*) neurons described by the equations:

$$\tau \frac{dG_i}{dt} = -G_i + \sum_{j=1}^{N} \omega_{ij} R_j$$
$$\tau \frac{dR_i}{dt} = -R_i + \frac{V_i}{1 + G_i}$$

where $i = 1, \ldots, N$ corresponds to individual alternatives, V_i is the value of the ith option, τ is an intrinsic timescale parameter, and the parameters ω_{ij} weight the input R_j to the gain neuron G_i. Using this dynamical normalization model, we examined the relationship between recurrent inhibition-mediated value normalization and the temporal evolution of value representation during the decision process. We also compared model predictions with existing and new observations of parietal neuron activity in the nonhuman primate.

Results: The dynamic model exhibits characteristic activity dynamics, with initial phasic transients preceding a steady-state level of activity. Specifically, dynamic normalization produces: 1) dynamic evolution to a stable fixed point, 2) normalized value coding at equilibrium, 3) time-varying value modulation during option evaluation, and 4) temporal asymmetry in the influence of option and contextual value information. These results are consistent with saccade-related parietal activity in both previous studies and new neural recordings. Notably, value coding is strongest during initial rather than late stages of valuation, suggesting that - in contrast to accumulator-based models of choice - economic decision-making may be more efficient at short timescales (a *speed-accuracy complementarity*).

Conclusions: These results show that a recurrent model of divisive normalization captures both the dynamic activity patterns and equilibrium value coding exhibited by cortical decision circuits. In addition to elucidating potential network mechanisms for decision-related neural activity, dynamic circuit models such as this one can predict novel patterns of choice activity outside the scope of traditional economic analysis.

Cross-modal coding in the anterior insula: shared and distinct neural codes for pain, disgust and unfair economic offers

A. Tusche,[1*] C. Corradi-Dell'Acqua,[2*] P. Vuilleumier,[2] and T. Singer.[1]

[1]Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig;
[2]Laboratory for Behavioral Neurology and Imaging of Cognition, University of Genève.

*Authors equally contributed to the study,
Correspondence: atusche@cbs.mpg.de

Objective: The anterior insula (AI) has been implicated in the neural encoding of a variety of cognitive and affective states, including the processing of painful, disgusting and unfair events. Evidence also suggests that this applies to events directed to oneself and when observing others exposed to these experiences. The present study aimed to investigate whether the overlap of relevant brain responses in the AI reflects domain-specific coding or a shared neural code for unpleasantness that is common to painful, disgusting and unfair economic experiences.

Methods: To address this question, we employed multivariate pattern analyses (MVPA) techniques that have previously been used to assess similarity between neural representations across experimental conditions. Using fMRI, we measured brain responses of 19 healthy participants while they performed 3 independent tasks. In each task, we presented aversive or neutral stimuli that were directed either to the participant or a confederate (resulting in a 2 [unpleasantness: aversive, neutral] x 2 [target: self, other] factorial design for each task). Tasks included a pain paradigm (noxious vs. non-noxious electrical stimulations), a taste paradigm (disgusting vs. neutral liquids) and an Ultimatum Game (UG) (unfair vs. moderately fair economic offers). Imaging data were analyzed using cross-classification techniques (MVPA) to test whether variance in neural response pattern evoked in the AI in one task was explained by response patterns in another task and modality.

Results: In line with previous evidence, we found that self-directed (participant) and vicarious experiences (confederate) elicited common neural signatures in the AI for the domain of pain, disgust and unfair economic offers (*cross-target* MVPA). Importantly, *cross-modal* MPVA also identified shared response patterns in AI across sensory modalities of painful and disgusting events, as well as response patterns specific to these basic sensory modalities. Finally, cross-modal analysis for the UG showed that response patterns elicited by unfair economic offers were similar to neural signatures for painful and disgusting events.

Conclusions: We found evidence for both domain-specific coding in the AI as well as for shared neural signatures across economic and basic sensory tasks coding for the unpleasantness of experiences. This finding suggests that the AI sub-serves multiple parallel processes during exposure to aversive events.

Understanding Neural Responses to Financial Offers using Facial Expressions

F. Rossi[1,*], V. van Son[2], I.R. Fasel[3], M.S. Bartlett[1,3] & A.G. Sanfey[2]

[1] Institute for Neural Computation, University of California, San Diego; [2] Donders Institute for Brain, Cognition and Behavior, Radboud University; [3] Emotient, Inc., San Diego.

[*]Correspondence at: frossi@ucsd.edu

Objective: It is challenging to measure and model the effects of emotions on economic decisions. We present a novel approach to address these issues. In particular, we seek further evidence for the notion that "Responders" in the Ultimatum Game (UG) experience moral disgust when receiving small financial offers **[1]**. We defined neural "fingerprints" of disgust by combining fMRI with a computer vision system for real-time analysis of facial expressions (Emotient SDK) **[2]**. Our goal was to test whether these fingerprints overlap with brain activity associated with small offers in the UG.

Methods: Twenty-six participants were scanned while playing a *Pictures Task* (PT) and the UG. In the PT, participants saw neutral, positive or disgusting IAPS images **[3]** (30 pictures × 2 runs). Additionally, participants played the role of "Responder" in the UG (25 offers × 2 runs). During the experiment, we collected videos of participants' faces, using the Emotient SDK to determine their facial display of disgust. We trained two classifiers, a "Voxels-to-Picture Classifier" (V2P) and a "Voxels-to-Face Classifier" (V2F), that used a distributed pattern of fMRI signal as input. The V2P predicts if the participant is observing a "disgusting" image, while V2F predicts the continuous measure of disgust assessed using facial expressions. Both models performed better than chance at predicting the picture types being observed.

Results: We tested whether the disgust detectors trained on PT data could predict small offers (€1-€3 out of €10) in the UG. While V2P performed poorly (accuracy = 52%; *ns*), V2F was very accurate (accuracy = 64%; $p < 0.0001$). The neural fingerprint used by V2F included voxels from anterior insula (R/L), amygdala (R), and caudate (L). V2F accuracy suggests that small offers in the UG trigger moral disgust, and that the rich description of affective experiences captured by facial expressions can be used to generate reliable multi-voxel patterns of emotions.

Conclusions: We were able to isolate a configuration of voxels associated with the facial expression of disgust in the PT, and we showed that this fingerprint predicts offer amount in the UG. Our findings reinforce the notion that small financial offers elicit moral disgust. Additionally, we provide a detailed neural representation of this experience, which generalizes from visual (PT) to moral (UG) disgust. The use of facial expressions to generate reliable neural fingerprints of emotions represents a substantial advance in the identification of the neural substrate of emotions, and may facilitate new tests of quantitative theories of the role of emotions in decision-making processes.

References:

[1] Sanfey, A., Rilling, J., Aronson, J., et al. (2003). The Neural Basis of Economic Decision-Making in the Ultimatum Game. *Science*, 300 (5626), pp. 1755–1758.

[2] Littlewort G., Whitehill, J., Wu, T., Fasel, I., et al. (2011). The Computer Expression Recognition Toolbox (CERT). *Proc. IEEE International Conference on Automatic Face and Gesture Recognition*, 17, pp. 298–305.

[3] Lang, P.J., Bradley, M.M., & Cuthbert, B.N. (2008). International affective picture system (IAPS): Affective ratings of pictures and instruction manual. *Technical Report A-8.* University of Florida, Gainesville, FL.

Propranolol selectively reduces loss aversion

Peter Sokol-Hessner[1*], Sandra F Lackovic[1], Russell H Tobe[2],
Bennett L Leventhal[2], Elizabeth A Phelps[1,2,3]

[1]Department of Psychology, New York University; [2]Nathan Kline Institute; [3]Center for Neural Science, New York University.

[*] Presenting author; psh234@nyu.edu

Character count: 2288 *(< 2300 characters, not incl. spaces)*

In risky monetary choice, recent correlational work has identified a role for emotions in the relative weighting of losses to gains. Individuals' unique loss/gain weighting, termed loss aversion, correlates with their relative physiological arousal responses to loss versus gain outcomes, suggesting that choices to avoid losses reflect individuals' experience of those losses. Additionally, amygdala hemodynamic responses to losses and gains show a similar pattern and correlation, and patients with amygdala lesions are less loss averse, suggesting the necessity of the amygdala's role. These prior findings support the hypothesis that amygdala-mediated arousal responses to loss and gain outcomes drive behavioral loss aversion, just as they mediate the effects of emotion on memory or avoidance behavior.

We sought to test this hypothesis by pharmacologically interfering directly with these responses using propranolol. Propranolol is a non-selective beta-adrenergic receptor antagonist which blunts arousal responses without sedative effects. As a highly lipophilic substance, propranolol crosses the blood-brain-barrier, and has been shown to eliminate the modulatory effects of arousal in other domains. Here, we administered propranolol (80mg) and placebo in a two-day, double-blind within-subjects design, hypothesizing that propranolol would reduce participants' loss aversion, which would provide evidence that the same adrenergic system supported the effect of arousal on risky monetary decision-making.

Participants completed an identical risky monetary choice task on each of two days. The choice task allowed us to quantitatively estimate and dissociate loss aversion, risk attitudes, and choice consistency for each participant and day. We found that propranolol reduced loss aversion, and the effect was both dose-dependent (a greater decrease for low BMI individuals) and baseline-dependent (a greater decrease for highly loss averse individuals). No other decision or valuation process was affected by propranolol.

Propranolol's reduction of loss aversion constitutes causal evidence that adrenergically-mediated arousal responses drive the relative weighting of monetary losses to gains, supporting the notion that emotions causally contribute to the computation of value, and converging with evidence connecting the adrenergic system to arousal's effects on memory and avoidance behaviors. By identifying this specific and causal relationship between a precise component of decision-making and its underlying neurohormonal system, we have here demonstrated a selective, causal, modulatory role of emotions in decision-making.

Acknowledgments:
This study was funded by the NIH, grant #AG039283

Contributions of Episodic Memory to Value-Based Decisions

K.D. Duncan,[1*] B.B. Doll[1,2], N.D. Daw[2], and D. Shohamy.[1]

[1]Department of Psychology, Columbia University; [2]Center for Neural Science and Department of Psychology, New York University.

[*]Correspondence at: kdd2110@columbia.edu

Objective: As a field, neuroeconomics has studied the influence of past outcomes on preferences by focusing on how abstracted values are incrementally learned over repeated experiences with a specific option. Many choices, however, are made between options with which we have limited past experience. These choices may be supported by the neurally distinct episodic memory system, which stores associative memories of distinct past experiences. To assess this possibility, we developed a new task in which value-based decisions can be made using both episodically and incrementally learned values.

Methods: Thirty participants played a computerized card game where on each of 250 trials they chose between two cards to win money. Participants could base their choice on two separate features: the color of the *deck* (binary, blue or red) or the picture of an *object* on the card (different for each card). The value of each deck was probabilistic and varied slowly over time, allowing participants to incrementally learn which deck had a greater probability of reward at any given point. Separately, each specific object perfectly predicted the card's value, allowing participants to use their memory of the outcome of a single previous trial to increase winnings. Critically, the value of the deck vs. the objects was decorrelated (r=.07) in our design by regularly reversing the decks' reward probabilities, allowing us to use computational models to estimate the extent to which each feature influenced participants' choices.

Results: In contrast to the assumption that value-based decisions are solely based on incremental learning, we found that participants' choices were best predicted by a model that included both incremental learning of deck values, estimated using a Q-learning algorithm, as well as one-shot episodic memory for object values. Moreover, the strength of episodic memories for distinct trials was modulated by the uncertainty of deck values: participants were more likely to remember and use object values from trials that closely followed a reversal in deck values. This effect was not driven by changes in reaction time during object-value learning and the reversal structure only influenced the *learning*, and not *use*, of object values, suggesting that participants did not simply attend more to objects following a reversal.

Conclusions: These results demonstrate that both incrementally learned values and memories for distinct episodes can drive value-based choices, and that the relative contributions of these memory systems may be modulated by uncertainty in the environment.

Acknowledgements:
This study was funded by the NIH/NINDS, grant number R01NS078784

A low-dimensional dynamical model of neural decision making

Jaron T. Colas[1]*, Alison M. Harris[3], & Antonio Rangel[1,2]

[1]Computation and Neural Systems Program & [2]Division of the Humanities and Social Sciences, California Institute of Technology
[3]Department of Psychology, Claremont McKenna College

*Correspondence: jcolas@caltech.edu

Objective: We introduce and test the simple competing accumulator (SCA) model, a connectionist model of a nonlinear dynamical system descended from the leaky, competing accumulator (LCA) model (Usher & McClelland, 2001).

Methods: The SCA model has been designed to be neurally plausible while balancing the constraint of minimizing its parameter count so as to ensure that each element is fully interpretable and avoid inappropriate assumptions and overfitting of empirical data. Furthermore, fitting the free parameters of a model of this complexity poses an intractably non-convex optimization problem with computational demands exacerbated by Monte Carlo simulation of stochastic time series lacking closed-form expressions. The SCA model is distinguished from the LCA model by core features with practical and theoretical implications, including greater overall simplicity and interpretability, reduction of less essential free parameters, stricter emphasis on lateral inhibition, and generation of explicit neural dynamics. The SCA framework is flexible and generalizes to decisions involving more than two alternatives with *n* attributes each while requiring only *(3 + n)* free parameters. Even when limited to behavioral data alone, the SCA model can emulate concurrent effects related to input magnitude and differences in magnitude that normative sequential-sampling models, such as the race model and the drift-diffusion model, qualitatively fail to capture.

Results: The SCA model was employed for three data sets acquired with different experimental paradigms incorporating electroencephalography and functional magnetic resonance imaging. The output of each instantiation of the model included accurate quantitative predictions of human subjects' choices and reaction times in addition to the corresponding temporal dynamics of aggregate neural activity in regions of the brain underlying value-based comparison processes. Correlates of these simulated neural signals were not only spatially localized with temporally coarse blood-oxygen-level-dependent signals, but also identified throughout high-resolution time courses of event-related potentials.

Conclusions: The SCA model stands as a viable and efficient tool that can provide a unified computational account of behavioral and neurophysiological data across diverse decision-making settings.

Changing exploration during value-based learning with frontopolar brain stimulation

Anjali Raja Beharelle,[1*] Rafael Polania[1], Todd Hare[1], and Christian Ruff[1]

[1]Lab for Social and Neural Systems Research, University of Zürich, Switzerland

[*]Correspondence to: anjali.rajabeharelle@econ.uzh.ch

Objective: Daily life involves frequent trade-offs between exploitation of known quantities and exploration of alternative opportunities. Furthermore, efficient exploration is critical for an organism's ability to seek out resources successfully. Frontopolar cortex (FPC) has been associated with decisions to explore rather than exploit and with tracking the reward probabilities of unchosen alternative options, however to date, evidence of FPC involvement in exploration is only based on correlations. Here we tested with brain stimulation whether FPC activity is causally necessary for exploratory decision making.

Methods: Seventy-nine healthy participants played a three-armed bandit task during a baseline period and while undergoing anodal, cathodal, or sham transcranial direct current stimulation (tDCS) of the right FPC. Choices were classified as exploratory or exploitative using a reinforcement learning model including parameters that quantified exploratory tendencies and the effects of relative bandit values on choice. All effects were quantified as changes relative to subjects' individual baseline behavior.

Results: Anodal stimulation increased exploratory behavior (i.e., subjects made more exploratory choices and showed larger shifts away from exploitative choices) while cathodal stimulation decreased exploratory behavior on those same measures. In addition, subjects receiving anodal stimulation earned less on the task, whereas subjects receiving cathodal stimulation earned more. Subjects with anodal stimulation also showed more exploratory tendencies based on model parameters. Sham stimulation had no effect on the task, and neither type of brain stimulation had unspecific cognitive effects on control tasks measuring ability to carry out calculations for the task, or risk aversion. Finally, subjects undergoing anodal stimulation underestimated the bandit rewards relative to their true value.

Conclusions: Right FPC is a critical neural component in driving exploration of alternative courses of action. Anodal stimulation of right FPC also causes subjects to underestimate the true value of the bandits, suggesting that exploration may be impelled by the notion that the grass is, in fact, greener on the other side.

Acknowledgements:
This study was funded by the Swiss National Science Foundation and the Marie Curie International Incoming Fellowship.

Title: A novel paradigm to evaluate the neural mechanisms of cooperation in non-human primates

Authors: Wei Song Ong (presenting author), Michael Platt
Institution: Duke University
Contact for corresponding author: weisong.ong@duke.edu

Deficits in social skills and communication define autism spectrum disorders. As social animals, a critical product of communication is the ability to cooperate--to work with another individual to achieve a common goal and compromise if needed. In brain areas where reward and value are represented, neuroimaging studies in humans have shown a signal associated with cooperative behavior. Single neuron recordings carried out in our lab have previously reported vicarious reward signals in the anterior cingulate gyrus and amygdala in monkeys when they work to deliver juice to a passive recipient monkey present in the same room (Chang et al., 2012). Here we describe a novel paradigm that requires the active engagement of both monkeys to achieve reward. We modified a variant of the hawk-dove economic game, which favors two mixed strategy equilibria. In the classic version of the game, the most profit is gained by switching strategies, thus keeping both players in constant competition.

In our task, 2 head restrained monkeys (M1 & M2) face each other above a shared LCD monitor placed parallel to the floor between them. Two colored annuli (1 green, 1 red) framing dot arrays and 4 targets (2 green, 2 red) are presented. Color indicates to which monkey the stimuli and targets belong. M1 & M2 indicate their choice of targets – one of which is larger than the other-- by controlling the motion of the dots within the annuli via joysticks. In 95% of the trials, the larger reward (denoted visually) is placed on the side of the screen distal to the controlling monkey, with the opponent's annulus interposed. On those trials, the smaller rewards are placed on the left of the screen. To achieve the larger reward, M1 should go straight. However, if the opponent (M2) also chooses to go straight for the larger reward on the same trials, the annuli collide and neither monkey receives a reward. If at least one monkey chooses to go left for the smaller reward, both receive their chosen rewards.

Four male monkeys consistently chose the larger reward in the absence of an opponent. They utilized the movement cued by the opponent's dot patch and avoided collision by picking the smaller reward on the left when needed. When we varied the coherence of the opponent's dots, thus varying cue certainty, the frequency of suboptimal choices increased as coherence decreased (threshold=13.2% coherence). In control trials where the larger reward was to the left, monkeys never make suboptimal choices regardless of the coherence of the opponent's dots. While 2 of the animals are actively engaged in this task, we will collect single neuron activity from dACC and mSTS, which is analogous to the human TPJ and shown to be active during cooperation.

Developmental Changes in Value-Based Learning

C.A. Hartley

[1]Sackler Institute, Weill Cornell Medical College

[*]Correspondence at: cah2031@med.cornell.edu

Objective: To date, the neuroscientific study of value-based learning and decision-making has focused primarily on their implementation in the adult brain. However, the neurocircuitry implicated in these processes undergoes substantial maturational changes during from childhood into adulthood, suggesting that value-based learning is likely to be qualitatively different in children and adolescents than in adults. Here, we present two studies examining how distinct reward-based learning processes change across development.

Methods: Studies of reinforcement learning distinguish two types of algorithms that can guide action selection. "Model-based" learning searches a cognitive model representing the consequences of potential actions, whereas "model-free" learning efficiently updates a cached action value associated with a state. In the first study, we had children, adolescents, and adults perform a two-stage reinforcement-learning task that can dissociate model-based and model-free contributions to choice in order to characterize developmental changes in the reliance on these two forms of learning

The value estimates that guide our actions can be acquired not only through our direct experience, but also through explicit instruction. In the second study, we examined whether the relative weighting of these forms of learning changes across development by placing instruction and experiential learning in competition in a probabilistic learning task.

Results: In study 1, we found that whereas the behavioral signature of model-free learning was present from childhood onwards, model-based influence on choice was not evident until adolescence, and continued to mature into adulthood. In study two, we found that whereas inaccurate instruction markedly biased adults' value estimation, children and adolescents placed greater weight upon their direct experience.

Conclusions: Collectively, these findings suggest that children and adolescents engage distinct neurocognitive processes as they learn to distinguish rewarding actions from those that are best avoided. For each study, we present a provisional model of how maturational changes in the brain might underlie these qualitative developmental changes in value-based learning.

Acknowledgements:
This work was funded by the Dewitt Wallace Readers Digest Fund and a generous gift from the Mortimer D. Sackler MD family.

Altered willingness to wait for delayed rewards in the context of psychopathology or focal brain injury

J.T. McGuire,[1*] D. Mukherjee,[1] R. Kazinka,[1] A.R. Vaidya,[2] and J.W. Kable[1]

[1]Department of Psychology, University of Pennsylvania
[2]Montreal Neurological Institute, Department of Neurology and Neurosurgery, McGill University

[*]Correspondence at: mcguirej@psych.upenn.edu

Objective: In uncertain environments, decision makers face a challenge in determining whether to sustain or curtail persistence toward delayed rewards. Previous work has demonstrated that willingness to wait is adaptively calibrated to the timing statistics of the relevant environment, and that activity in ventromedial prefrontal cortex (VMPFC) tracks the value of future outcomes in a dynamic and context-sensitive manner. Here we examined how persistence might be altered by VMPFC-specific brain injury or by clinical depression, a psychiatric disorder linked to VMPFC dysfunction.

Methods: Experiment 1 compared a group of individuals diagnosed with Major Depressive Disorder (MDD) to a demographically matched control group. Experiment 2 compared groups of individuals with focal brain lesions in VMPFC or lateral prefrontal cortex (LPFC) and a demographically matched control group. All participants performed a foraging-like task that involved voluntary waiting for randomly timed rewards. The initial task environment was constructed so that high persistence was the ideal strategy. Participants subsequently encountered a second environment that favored more limited persistence. We used survival analysis methods to quantify the average amount of time each individual was willing to wait for delayed rewards in each environment.

Results: Analyses of group differences focused on the task environment in which high persistence was the normative strategy (results suggested that behavior in Block 2 of the experiment was heavily influenced by experience during Block 1). Experiment 1 showed significantly reduced persistence in the MDD group relative to the control group. Experiment 2 showed a trend toward a large but variable reduction in persistence in the VMPFC-lesion group relative to the other groups. In each experiment, analysis of performance over time suggested the greatest group differences emerged late in the task block rather than merely reflecting slower initial learning.

Conclusions: Results suggest that both psychopathology and VMPFC injury may impair decision makers' ability to adapt to an environment in which normative behavior entails high willingness to wait for delayed outcomes. These results are compatible with previous fMRI findings implicating VMPFC activation dynamics in successful persistence behavior. A key open question is whether the observed effects emerge from a bias toward low persistence (in which case the MDD and VMPFC groups might perform successfully if initially placed in an environment in which limited persistence was adaptive) or, alternatively, from a general deficit in context-sensitive calibration of behavior.

Acknowledgements:
This work was supported by NIH grants R01-MH098899 and R01-DA029149 to JWK, and F32-DA030870 to JTM.

Selection bias for nearer objects is supported by neural signals in the nucleus accumbens

Sara E. Morrison[1*] and Saleem M. Nicola[1,2]

[1]Dept. of Psychiatry, [2]Dept. of Neuroscience, Albert Einstein College of Medicine

*Correspondence at sara.morrison@einstein.yu.edu

Objective: People and animals show a pronounced preference for rewarding options that are nearby as opposed to farther away, a phenomenon we call "proximate reward bias." Although this phenomenon has long been recognized by psychologists, it has received scant attention from neuroscientists, and the neural mechanisms by which proximity biases decision-making remain unclear. We hypothesized that proximity-modulated signals in the nucleus accumbens (NAc) might bias animals to choose nearer objects. This could occur in two ways: via cognitive evaluation of the options, or via a non-cognitive process by which nearer objects more strongly elicit conditioned approach.

Methods: To distinguish between these alternatives, we designed a task in which proximity to the choice targets (levers to press) was highly variable, then systematically varied the reward size and effort cost associated with each target. While rats performed this task, we recorded the activity of individual neurons in the NAc.

Results: We found that proximity was a major determinant of suboptimal choices, and that subjects often chose the nearer option even when it resulted in greater effort expenditure and delay to reward. Therefore, proximate reward bias was unlikely to be caused by effort or delay discounting. Cue-evoked activity in the NAc robustly encoded proximity to a lever, regardless of which lever – optimal or suboptimal – was subsequently chosen. The same activity showed no overall preference for high reward or low effort, and did not track the value of individual choice targets. Surprisingly, there was no correlation between the representations of reward size and effort level within single neurons; nor was proximity encoded jointly with reward size or effort. Thus, it appears that the NAc does not encode the value of stimuli in a common currency.

Conclusions: Cue-evoked signals in the NAc are likely to underlie proximate reward bias via a non-cognitive mechanism: promoting impulsive approach to the nearest reward-associated choice target. In contrast, NAc activity is ill suited for supporting choice based on anticipated reward or effort. We note that the most frequently studied decision variables – e.g., expected reward, effort, and delay to reward – paint an incomplete picture of choice behavior: proximity to reward is encoded independently and influences decisions on its own.

Acknowledgements: Funding provided by NIH grants to S.M.N. (DA019473, DA038412, MH092757) and S.E.M. (DA034465); NARSAD, the Klarman Family Foundation, and the Peter F. McManus Charitable Trust to S.M.N.; and the Charles H. Revson Foundation to S.E.M.

Cumulative evidence that dopamine-associated cached values are not used in action selection

Nick G. Hollon[1,2,3], Monica M. Arnold[1,2], Vicente Martinez[1,2], Mark E. Walton[4] and Paul E. M. Phillips[1,2,3*]

[1]Department of Psychiatry & Behavioral Sciences, [2]Department of Pharmacology and [3]Program in Neurobiology & Behavior, University of Washington. [4]Department of Experimental Psychology, University of Oxford.

[*]Correspondence at: pemp@uw.edu

Objective: In the foundations of the field of neuroeconomics is the notion that dopamine signaling updates cached values associated with stimuli or actions and that these values are used to compare available options in making choices. Here we critically evaluated the relationship between dopamine-associated cached values and action selection.

Methods: Dopamine release was monitored in the nucleus accumbens of rats performing decision making tasks. We presented stimuli associated with different actions to the rats unexpectedly to generate dopamine-mediated prediction-error signals that would provide a read out of the cached value. We measures the cached values during devaluation of rewards, during choices requiring the tradeoff between costs and benefits and under conditions that promoted behavioral inflexibility.

Results: Under each of the behavioral tasks we were observed conditions where animals would consistently choose an option that was not associated with the largest cached value. Following reward devaluation, rats' choices reflected the devalued reward immediately but the cached values did not initially and updated following experiential pairing of the cue with the reward at its new value. In tradeoffs between reward size and response requirement, the cached value robustly incorporated the benefit, but not the cost of the option and thus, did not account for all of the economic parameters used in the choice. Indeed, by manipulating the balance between the response cost and the reward size, we could establish conditions where the cached value was reliably highest for the non-preferred option. This finding was replicated even under conditions of behavioral inflexibility (c.f., habit) where model-free valuation systems are thought to predominate. Moreover, systemic dopamine antagonists increased the number of omitted trials, but did not change the allocation of choices between options.

Conclusions: These data indicate that dopamine-associated cached values are not used to decide which of the available options to select, but rather they to energize an action once it has been selected.

Acknowledgements:
This study was funded by the NIH, grants R01-MH079292, R01-DA027858 and R01-AG044839.