

Decision making: Neural prediction errors show risk sensitivity

Yael Niv¹, Peter Dayan², and John P. O’Doherty³

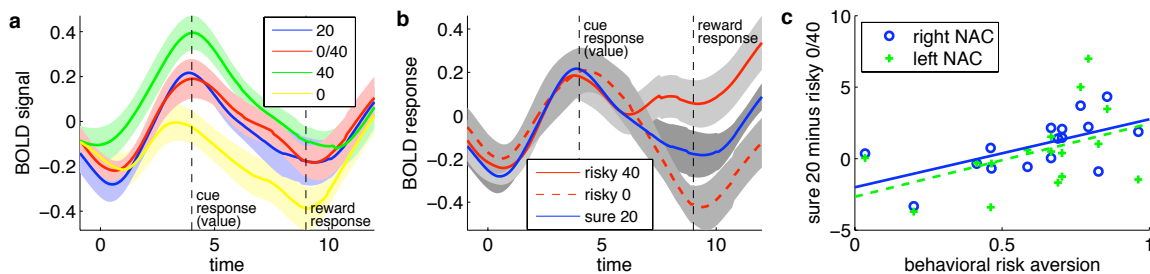
¹Psychology Department, Princeton University, ²Gatsby Computational Neuroscience Unit, UCL

³Division of Humanities & Social Sciences and Computation & Neural Systems Program, Caltech

Decision making is influenced not only by the expected reward value of options, but also by their variance or risks. Traditional reinforcement learning models of action selection, however, use temporal difference (TD) methods to learn the *mean* value of an option, ignoring risk. These models have been strongly linked to learning via prediction errors conveyed by dopaminergic neurons. Although variance clearly influences choice behavior, risk sensitivity does not have to be manifest at the level of TD learning: the dopaminergic system may be learning expected values, while risk is monitored in other systems, with both contributing to the final decision. Here, we used fMRI in a human decision-making task, in order to determine whether correlates of risk sensitivity can be seen in neural prediction error signals themselves. Our results reveal that choice behavior is better accounted for by incorporating risk-sensitivity into TD learning, and, furthermore, that the BOLD correlates of prediction error learning in the brain indeed reflect subjective risk-sensitivity.

Humans performed a choice task in which three options were rewarded with 0, 20 and 40 cents, respectively, and one option was rewarded with 0 or 40 cents with equal probabilities. Some trials involved choosing between the ‘sure’ 20 cents option and the ‘risky’ 0/40 cents option. Preference for the ‘sure’ option was used to quantify each subject’s degree of risk-aversion. Although the traditional TD model can explain risk sensitivity as a result of choice-induced sampling biases, model comparison showed that a recently proposed risk-sensitive TD model predicted subjects’ behavioral choices significantly better.

A critical difference between the models is that the traditional TD model predicts that the ‘risky’ option will, on average, have the same value as the ‘sure’ option, while risk-sensitive TD penalizes variance, making the value of the ‘risky’ option lower than that of the ‘sure’ option for risk-averse subjects, and vice versa for risk-seeking subjects. To differentiate between the models on the level of neural TD learning, we therefore assessed the correlation between behavioral risk-preference and the valuation of the ‘sure’ and ‘risky’ options derived from the BOLD signal in anatomically defined regions of interest in the left and right nucleus accumbens (NAC) of each subject. First we confirmed, for the first time using a model-free non-exploratory analysis, that both right and left NAC BOLD signals correspond to a prediction error signal (Figure 1a,b), and, at least in our task, do not encode for other quantities of interest (such as reward magnitude). We then used the BOLD prediction error signal at trial onsets to derive the subjective mean value of each option. A regression of the behavioral risk preferences and the subject-by-subject difference between the values of the ‘safe’ option and the ‘risky’ option showed a significant positive correlation in both the left and the right NAC (Figure 1c). This suggests that risk-sensitivity can be seen at the level of TD learning in the brain.



Acknowledgments. We are indebted to Jeffrey Edlund for acquiring the imaging data, and to Laura deSouza for marking the anatomical ROIs. Supported by a Human Frontiers Science Program fellowship to YN.