# Computational, Cognitive, and Neural Models of Decision-making Biases

**Jonathan Malmaud (malmaud@mit.edu), Joshua B. Tenenbaum (jbt@mit.edu)**
Brain and Cognitive Sciences, 43 Vassar St.
Cambridge, MA 02139 USA

**Peter Dayan (dayan@gatsby.ucl.ac.uk)**
Gatsby Computational Neuroscience Unit
17 Queen Square
London, WCIN 3AR England

**Laurence T. Maloney (laurence.maloney@nyu.edu)**
Dept. of Psychology, 6 Washington Place
New York, NY 10030 USA

**Edward Vul (evul@ucsd.edu)**
Dept. of Psychology, 9500 Gilman Dr. 0109
La Jolla, CA 92093 USA

**Nick Chater (n.chater@ucl.ac.uk)**
Division of Psychology and Language Sciences
26 Bedform Way
London, WC1H 0AP England

## Summary

The question for the symposium is how best to understand biases in decision-making, going beyond traditional judgment and decision-making (JDM) accounts such as prospect theory to take a more modern reverse-engineering perspective bridging rational computational, algorithmic, and neural levels of explanation, and viewing decision-making under risk and uncertainty not just as a simple matter of evaluating lotteries but in the context of cognition more broadly, taking seriously learning, perception, motor control, memory, and action planning.

The dominant normative approach to studying decision-making under risk is axiomatic expected utility theory, which argues that any agent obeying seemingly reasonable axioms of choice consistency can be modeled as maximizing the expected utility of its decisions. From decades of research that analyzes people's choices between simple gambles in the lab, it is known that humans routinely violate these axioms. This has forced decision theorists to adopt descriptive models of choice that lack a normative rational in order to account for observed patterns of choice, the most prominent of which is Khaneman's and Tversky's prospect theory for one-shot decisions under risk with immediate outcomes and hyperbolic discounting for decisions involving delayed outcomes.

There are several challenges not addressed by prospect theory and its variants. First, they are silent on the issue of the cognitive mechanisms that are actually responsible for human choice behavior. Second, they do not seem as a practical matter to scale to real-world decision problems, where the space of possible outcomes and actions is not sharply defined, the effects of actions are highly uncertain, and the explicit calculation of expected values is impractical. Third, they do not strongly constrain or give an underlying rationale for the probability weighting function, temporal discounting function, or utility function featured in prospect theory. Thus these models cannot explain why these functions' estimated forms and parameters seem to be greatly affected by seemingly irrelevant factors of the task framing and setup, such as whether the outcome probabilities are presented numerically in tables or learned through experience and why the evaluation of individual gambles seems to be highly effected by the properties of other gambles in the choice set. More broadly, these theories fail to explain why in day-to-day life human decision-making seems to generally be highly robust and effective while sharply contrasting with normative predictions in the simple, stylized decision tasks commonly used in JDM experiments.

This symposium brings together researchers who represent a variety of perspectives on ways cognitive science can inform our understanding of decision biases to address these challenges, with relevance at all three Marr levels of analysis. Malmaud and Tenenbaum, and Dayan both offer computational-level Bayesian accounts that explain decision-making biases as resulting from reasoning with priors that are adapted for real-world or evolutionary-relevant decision tasks. Malmaud and Tenenbaum explain choices in terms of advanced models from the AI planning literature and animal foraging theory. Dayan offers a neurobiological implementation of inference that spans the Marr levels.

Other approaches relate to algorithms levels of the Marr hierarchy with links to lower and higher levels. Vul offers an algorithmic description of biases as resulting from cognitive limitations associated with reasoning using only a limited number of samples from a posterior over decision parameters. Maloney and Chater link high-level decision-making to known properties of perception and cognition, such as scale-invariance. Maloney gives a unifying account of the probability weighting function as arising from the same principles as perception of continuous quantities in psychophysics. Chater explores the origin of subjective utility and temporal discounting through connections to broader cognitive processes.

One general idea that cuts across all these approaches is that human decision-making can be modeled in a unified way as the result of general cognitive principles that offer prin-

cipled explanatory accounts of biases in decision-making, rather than via a series of descriptive utility-maximizing models that have undergone ad hod adjustments to account for a mélange of deviations from a narrow normative standard.

## Malmaud and Tenenbaum: Prospect theory as rational response

We will open the symposium by presenting a brief review of the traditions of axiomatic decision theory and descriptive prospect theory, including how sophisticated computational models are beginning to fill in some of prospect theory's known shortcomings as a model of high-level decision-making at the individual level. We will briefly discuss our position that human decision-making is adapted for solving rich, sequential decision problems with structured goals and highly uncertain action-outcome contingencies and as such should not be expected to perform optimally according to narrow normative standards in simple one-shot decision tasks with known contingencies. We will show how modeling human choices as the result of employing state-of-the-art AI methods for planning under uncertainty to a specific class of 'survival' goals commonly studied in the animal foraging literature naturally implies a sequential decision strategy that is compatible with the descriptive predictions of prospect theory. Our approach is also able to make predictions about human behavior for a wide class of tasks for which prospect theory is not applicable. We will present preliminary empirical evidence that these predictions are supported on a specific set of ecologically relevant sequential decision tasks.

## Dayan: Pavlovian choice illusions

One useful interpretation of many perceptual illusions is in terms of biases resulting when the mechanisms of inference reflect genetically encoded or learned priors that are inconsistent with a given scene. We will consider how some of the biases of decision-making on which this symposium focuses can be seen as illusions of choice arising from forms of Pavlovian influences that reflect evolutionarily appropriate dispositions. These influences are exerted by various neural systems, notably the neuromodulators dopamine and serotonin and regions of the amygdala, acting on areas such as the striatum that are involved in decision-making. As an example, we will discuss the case of behavioral inhibition, which is one very general response to potential threats, and is closely associated with serotonin. We will show how such inhibition can lead to a particular form of bias in the on-line evaluation of complex options, and show how problems with this bias might even have deleterious psychiatric consequences.

## Maloney: Ubiquitous log odds

Similar patterns of distortion are found in visual frequency estimation, frequency estimation based on memory, and in the use of probability in decision-making under risk. Based on joint work with Zhang et al., I will show that probability distortions in all cases (so far) can be approximated by a linear transformation of the log-odds of probability or relative frequency. The slope and intercept of the linear transformation control probability distortion. Researchers have not been able to predict or explain the values of slope and intercept observed in experiments across tasks or across participants.

In Zhang & Maloney (2012) we focused on one method for presenting probability, the relative frequency of items of one kind in a visual array of N items. We developed a model of human distortion of relative frequency and demonstrated in two experiments that we can separately control slope and intercept with high accuracy. Our results support the conjecture that probability is systematically adapted to particular tasks much as perceptual information concerning lightness or loudness is transformed. We shown how a simple model based on chunking of information can explain the results we observe with a high degree of precision.

## Vul: Decision biases and heuristics arising from inference by sampling

Across many domains, people integrate sophisticated world knowledge with prior expectations nearly optimally, yet when making conscious cognitive judgments, they seem to be grossly irrational. I will explore a potential explanation: that conscious cognitive judgments reflect sample-based approximate inference under constrained cognitive resources. Experiments measuring multiple judgments from individuals with no new information yield evidence for this sampling proposal: any one decision appears to reflect only a small fraction of the information the participant has available, suggesting that each decision is based on only a small number of samples. Here, I will talk about the tradeoffs inherent in using a small number of samples for a decision: why we might want to use few samples, the consequences of using a few samples for judgments, the risks associated with using a few samples when rewards are asymmetric, and how these consequences relate to biases seen in judgment and decision-making.

## Chater: From cognitive principles to JDM

This talk will consider how far candidate cognitive principles (such as scale-invariance, relative coding of magnitudes, and incommensurability between distinct dimensions) can provide quantitative and qualitative explanations of results in decision-making. I will illustrate how widespread patterns in JDM (such as constant relative risk-aversion and hyperbolic time-discounting) can be derived; and consider how basic cognitive processes can explain when and in what way, such regularities break down. The aim is to build a theory of JDM built on cognitive principles, rather than rational axiomatic foundations.