

The Neurobiological Foundations of Valuation in Human Decision Making under Uncertainty

Peter Bossaerts¹ and Ming Hsu²

1. Introduction³

The goal of this chapter is to review recent neurobiological evidence to improve our understanding of human valuation under uncertainty. Although we are ultimately interested in human behavior, we will borrow from studies of animals with related brain structures, namely, non-human primates. Specifically, we wish to explore how valuation is accomplished? As we shall see, the evidence rejects a pure “retrieval from memory” model. Instead, values are *computed*. This raises the issue: what computational model(s) are being used? Specifically, since actual choice can be summarized in terms of a single-dimensional utility index such as expected utility or prospect theory, we are effectively wondering how such an index is computed. We wish to understand the effect of perceptual biases on this computation, as well as the role of emotions. How does the computational model generate the risk aversion that we may see in choices? Or, in ambiguous situations, how is ambiguity aversion revealed in choices (Hsu, Bhatt et al. 2005; Bali, Demirtas et al. 2006; Huettel, Stowe et al. 2006), and what model underlies it, e.g., alpha-maxmin preferences (Ghirardato, Maccheroni et al. 2004), anticipated regret (Segal 1987), or some other.

A natural question for economists is, if choice can be represented “as if” some utility index is maximized, why should we bother studying the computational aspects behind choice. Our quest for the computational model underlying valuation has two purposes. First, we hope to be able to improve choice prediction. We consider the utility index model as a reduced-form model of choice, which, in analogy with reduced-form models of the macro economy are sensitive to circumstances, i.e., whose parameters need to be revised every time one moves from one framework to another. This is the so-called Lucas critique (Lucas 1976). This sensitivity, or lack of robustness, has been frequently noted in the literature as “framing effects”. Our computational model, in contrast, is meant to be structural, and

¹ Ecole Polytechnique Fédérale Lausanne (EPFL), California Institute of Technology (Caltech) and Centre for Economic Policy Research (CEPR)

² Beckman Institute for Advanced Science and Technology, and Department of Economics, University of Illinois at Urbana-Champaign

³ Financial support from the Swiss Finance Institute and the Swiss National Science Foundation (Finrisk NCCR) is gratefully acknowledged.

hence, has a chance of outperforming the reduced-form utility index model. Of course, the proof is the pudding.

Not only should a computational model enable us to accommodate framing biases, it should be capable of capturing part, if not all, of the randomness one usually adds to a utility-based choice model to accommodate the erratic nature of observed choice. The traditional approach in applied economics has been to consider the utility index itself to be random, and to attribute certain distributional properties to this randomness that ultimately leads to a choice model based on logistic regression (McFadden 1974). The goal of exploring computational models of choice is to provide a different (computational) basis for the erratic nature of observed choice.

Artificial intelligence is the second reason why we are interested in the actual neurobiological computational model behind observed choices. Modern decision theory has been quite successful in prescribing learning and choice in well-defined situations. Once the complexity of the problem increases, however, decision theory has little concrete to propose. One case in point is the Iowa Gambling Task (IGT) (Bechara and Damasio 2005), which is essentially a complex four-armed bandit problem. To date, no-one has been able to spell out the optimal strategy in this task. Humans, however, are quite capable of quickly learning the ordering (in terms of reward and risk) of the four bandits (Bechara, Damasio et al. 1997).

However, the learning model behind human choice in difficult situations such as the IGT is almost surely not Bayesian, despite the popularity of this approach. Open-ended situations almost by definition involve high-dimensional parameter spaces, and it is well known that Bayesian learning generically fails to converge when the parameter space becomes large, while simple learning algorithms based on classical statistics often do a very good job in capturing the essence of the stochasticity at hand (Diaconis and Freedman 1986). This in turn leaves us the question of how do humans learn in such situations? What computational model are they using? If we knew this model, perhaps we could enrich decision theory to make it work in more open-ended or even ill-defined decision situations.

One clue in the search is the plethora of brain lesion data, much of it using the IGT, supporting the idea that specific brain regions underpin these computations. In particular, this involves the ventromedial prefrontal cortex and the amygdala, as patients with lesions to these brain regions are impaired in the IGT while at the same time show normal or even superior performance in working memory, IQ, and decisions and games that are less open-ended (Bechara, Damasio et al. 1997) (Bechara, Damasio et al. 1999).

In this chapter, we focus on valuation in the context of choice under uncertainty. The computational model that emerges from a study of the neurobiological data is one where expected rewards are traded off against risk (or, in the case of risk-seeking behavior, risk is added to expected reward). This model has been popular in finance, where risk is usually represented in terms of reward variance (although more advanced models take into account higher moments); whence its name: the mean-variance model.

The link between the utility models representing actual choice under uncertainty, such as expected utility and prospect theory, on the one hand, and the mean-risk model, on the other hand, is often clarified by means of Taylor series expansions (Bali, Demirtas et al. 2006). One of the goals of this chapter is to demonstrate that the logic of a computational model based on a trade-off between expected reward and risk can be extended to choice under ambiguity as well. More specifically, choice under α -maxmin preferences can be computed from trading off expected reward against ambiguity level.

Throughout this chapter, we take the position that valuation is an exercise that stands on its own. That is, valuation makes sense even in the absence of choice. It makes sense because computing values take energy, and hence, an organism that has learned to correctly compute the values of options it is forced to take on (imperative trials) will also be better able to determine the adequate choices when there is freedom (free-choice trials). As we shall see, brain activation when a subject is faced with gambles with uncertain outcomes is comparable across imperative and free-choice trials.

When valuation is a pursuit distinct from choice, values revealed through choice may be different from the valuations that come out of the (choice-independent) computations. One then wonders whether we, as economists, care at all to know what the latter are. After all, economists are only interested in choice. We shall spend some time elaborating on this issue. Suffice it to say here that the distinction is important for welfare assessment.

The remainder of this chapter is organized as follows. The next section elaborates on the distinction between values the brain computes and the values revealed through choice. Section 3 reviews the neurobiological foundations of valuation under pure risk. Section 4 extends the logic of expected reward and risk to situations with ambiguity. Note that we will not explore the neurobiological foundations of the translation of valuation into choice. The reader will find ample coverage of this topic in later chapters in this volume.

2. Distinguishing Valuation and Choice

In many human endeavors, valuation is performed even in the absence of any immediate necessity to choose. Again, finance is a case in point, in part because financial valuation is often complex and time-consuming, while good choice opportunities are rare and short-lived. The cost of computing values provides a normative rationale for why valuation may be done in the absence of free choice. Figure 1, adapted from (Berns, Capra et al. 2007) provides strong evidence that valuation is performed even in imperative trials, and in a way that is relevant for free-choice trials. It shows how well a valuation model fit to neuronal activation during imperative trials can predict choice in free-choice trials. In fact, the neurobiological choice model predicts choices better than a utility-index-based model estimated from the choices themselves! This demonstrates not only that valuation is done during imperative trials, but that the resulting values are relevant for choice in free-choice trials as well.

Although brain activation during imperative trials reflects valuations that are compatible with the values revealed in free-choice trials, and hence, that brain activation in imperative trials can be used to predict choice in free-choice trials (Figure 1), the fit is not 100%. In essence, the neurobiological data suggest that there are (at least) two value signals: one revealed through activation in brain regions not directly involved in the physical implementation of choice; and a second one revealed through activation of the neurons controlling the physical act of choice. Given their interest in choice, one then wonders to what extent economists should care to know the former value signal.

The most important argument concerns welfare. What is it that the subject really wanted? Indeed, the mere presence of two (or more) valuation signals suggests that there may be a difference between what subjects wanted and what they did. Of course, we may probably never be able to determine what exactly subjects really wanted (i.e., to extract their true preferences), but a better understanding of how the various value signals are put together may help us better appreciate the difference between revealed and true preferences.

It is fundamental to question coincidence of revealed and true preferences. Is a utility-index model such as prospect theory a preference theory? Prospect theory is known to fit choice well, but do the values of its utility index reflect real preferences? The insistence on deriving choice-theoretic axiomatic foundations for prospect theory (Kahneman and Tversky 1992) indicates that economists consider prospect theory to be a preference theory. The risk-seeking attitude towards losses that is implied by prospect theory, for instance, may not reflect intended behavior, but a bias imposed by some kind of “default action” that is deeply engrained in organisms that try to maximize the probability of survival in the face of adverse events.

Recent neurobiological evidence supports the latter interpretation. Choice in loss situations appears to be associated with activation in amygdala, irrespective of a subject’s tendency to take on risk. Orbitofrontal cortex (OFC) activation, however, decreases with subjects’ tendency to become risk loving for losses (Figure 2). These data suggest a very different valuation model underlying choice, based on “default actions” that need to be overcome in situations where these actions are inappropriate – yet not everyone manages to over-ride the default. In this interpretation of the data, amygdala activation engages the “default action”, while OFC is engaged in over-riding this choice when inappropriate. Perhaps no one wants to be risk-seeking for monetary losses, but many subjects cannot overcome their tendency to become risk-seeking when conditions are adverse.

More generally, it is not so much the value signals that may be discerned in, e.g., prefrontal cortex (vmPFC), which are of importance. More relevant is the way the signals are computed: What brain regions are involved? What is the nature of their signals? How do they correlate with behavior?

The potential difference between true and revealed preferences is in principle not important for predicting choice; yet it is of utmost important for welfare. At a mundane level, one should ask whether bankers and brokers ought to engage in eliciting the right prospect-theoretic utility index from their customers, through, e.g., questionnaires on past or hypothetical choices, to be used to determine optimal portfolios. Only if prospect theory captures the true preferences does this make sense. If it does

not, one wonders what the point is of deriving optimal portfolio choice for prospect-theoretic preference profiles (Gomes 2005).

In fact, the mere difficulty of computing optimal choices should cast doubt on the interpretation of prospect theory as a preference theory. Indeed, an organism that truly behaves in a prospect-theoretic way must be endowed with tremendous capacity to compute optimal choices. To put it more bluntly: perhaps organisms would be better off with logarithmic preferences, for which choices are easily computed, e.g., one can be myopic (Hakansson 1971). In addition, it has the advantage that choices maximize survival probabilities (Blume and Easley 1992)!

De Martino e.a.'s data suggested that choice is the result of valuing the appropriateness of some "default action" given the stimuli at hand – and, if inappropriate, effort to overcome one's tendency to implement the default action. Recent single-unit recording of neurons in the monkey brain indicates that caudate plays a crucial role in valuation of default actions. Figure 3, adapted from (Lauwereyns, Takikawa et al. 2002), provide evidence.

In the task that generated Figure 3, the monkey was asked to fixate its eyes in the middle of the screen, and to move its eyes (execute a "saccade") at a certain point in the direction indicated by a prior stimulus. Only one direction was rewarded, however. Actually, to be precise, the chance that one direction was rewarded was an order of magnitude higher than the other direction. As it turns out, certain caudate neurons increase their firing rate at trial onset, even before the stimulus, and hence, the required direction of motion of the eyes as well as the amount of the reward, is known. This firing reflects valuation of the default action, namely, to move the eyes in the direction that is most likely to be rewarded. When the stimulus appears, neuronal firing either increases, reflecting confirmation that the default action is correct, or decreases, signaling that the default action is inappropriate (Figure 3). In fact, if percentage correct saccades are an indication, effort is involved in avoiding the default action. Indeed, Lauwereyns et al. (2002) reports that the monkey makes more mistakes, i.e., moves its eyes in the wrong direction, when the stimulus requires a saccade in the less rewarding direction.

The firing pattern of the caudate neurons also casts doubt on random utility models (McFadden 1974) as an explanation of behavior that is at odds with maximization of a fixed utility index. Instead, it suggests that the apparent "erratic" behavior of the monkey (its mistakes) is the result of its inability to overcome the default action. E.g., if the monkey moved its eyes in the direction that was rewarded with higher probability while the stimulus instructed it otherwise, it was not that its risk tolerance had increased so much that it wanted to take the chance that its visual inference (of the instructions revealed by the stimulus) was wrong. The mistake was caused by the monkey's inability to overcome its default action (a saccade in the direction of reward with higher probability). Conversely, if the monkey did not make a mistake and moved its eyes in the direction of lower probability of reward, as instructed through the stimulus, it is not that the monkey became less risk tolerant and did not want to take chances. Instead, it overcame its natural inclination to saccade in the direction that tended to be rewarded with higher probability.

Even if the random utility model and the “default action” valuation model predict choice equally well, economists should be aware that the neurobiological evidence favors the latter. This is because the welfare implications of the two models are diametrically opposed. According to the random utility model, choice is always optimal. The agent makes different choices in two instances of the same situation only because utility fluctuates. In the “default action” valuation model, however, choice is often sub-optimal: the monkey wanted to move its eyes in the opposite direction but could not because of the effort to overcome its habit. If some external correcting device can be provided at an effort that is less than that required for the monkey to personally overcome its action bias, welfare improves. See Bernheim and Rangel (2006) for further discussion of welfare assessment in a context where one needs to distinguish between “true” and “decision” utilities.

3. Valuation under pure risk: trading off risk against reward

Pure risk is a situation where probabilities are known, either because they are given, or because the agent has gone through lengthy training. Economists model choice under pure risk as if the agent maximizes a utility index based on a separation of probabilities (of all possible states of nature) and utilities (of rewards in a state). This separation is known as *probabilistic sophistication* (Marinacci 2002), and is common to virtually all preference models under pure risk (including prospect theory).

There appears, however, to be little neurobiological evidence for value computation based on separation of (state) probabilities and (reward) utilities. Instead, brain activation data suggest a separation of encoding of expected reward and of risk.

Evidence of encoding of expected reward, and the crucial role of the dopaminergic system goes back a long way, to investigation of firing of dopaminergic neurons in the monkey midbrain in the face of stochastic rewards (Hollerman and Schultz 1998). Closer study of the activation of these neurons has revealed that the brain uses a simple yet versatile and powerful algorithm to learn expected rewards based on the Rescorla-Wagner reinforcement learning rule and extendible to complex multiple stimuli-reward situations (Montague, Dayan et al. 1996) (McClure, Berns et al. 2003) (O’Doherty, Dayan et al. 2003). At the core of this algorithm is a prediction error – the difference between actual and expected (sums of discounted) rewards. A number of chapters in this volume provide details, so we shall not elaborate here.

One crucial element is missing in this account of (expected) reward learning, namely, risk. Specifically, to learn at the optimal rate, it is important that the decision maker assesses the risk of making a prediction error. The intuition is very simple. If one expects to make large prediction errors, i.e., if the risk is

expected to be high, then one should not change one's prediction much based on a large actual prediction error. In other words, the learning rate should be low⁴.

There are (at least) two ways in which risk assessment could affect learning. One is directly through the learning rate; another one is through adjustment of the prediction error. Indeed, a simple scaling of prediction errors by their expected size accomplishes the same. See (Preuschoff and Bossaerts 2007) for details. Evidence for such a type of "adaptive encoding" was recently discovered in firing of dopamine neurons in the monkey brain. See Figure 4: shown are average firing rates and single-trial firings of dopamine neurons for three types of trials, in the order (top to bottom panel) of increasing expected size of the prediction error. When reward occurs, the prediction errors are insensitive to the level of anticipated risk, demonstrating that they are somehow "scaled."

(Expected) reward learning is complex. Not only should it be affected by the expected size of the prediction error (risk), other considerations play a role, such as how much an optimal predictor should co-vary with the prediction error. For instance, in situations where the environment changes rapidly, past prediction errors become obsolete fast, and hence prediction should rely more on recent prediction errors. Effectively, this means that the learning rate should increase. The intuition has a rigorous underpinning [see (Preuschoff and Bossaerts 2007)]. (Behrens, Woolrich et al. 2007) recently provided evidence that humans do adjust their learning rate to the degree of stability. We shall not elaborate here.

In addition to being necessary for optimal learning, risk encoding could also play a role in the decision-maker's risk *attitude*. Risk encoding is actually *presumed* to occur in risk-sensitive organisms. How else would their behavior be affected by risk? Not surprisingly, by now the evidence of risk encoding in the human and non-human primate brain is overwhelming. Regions where activation appears to be sensitive to risk include insula, anterior cingulate cortex (ACC) and inferior frontal gyrus (IFG) [e.g., (Critchley, Mathias et al. 2001), (Paulus, Rogalsky et al. 2003), (Huettel, Song et al. 2005)]. Some of these regions, such as insula, seem to encode risk *exclusively* (Preuschoff, Bossaerts et al. 2006) (Preuschoff, Quartz et al. 2007).

Moreover, in simple monetary gambles, activation seems to reflect *variance* (or its square root, standard deviation), i.e., the expectation of the squared prediction errors (Preuschoff, Bossaerts et al. 2006) (Preuschoff, Quartz et al. 2007). As such, the encoding provides a key parameter with which to compute the utility of a risk-averse agent. Indeed, a Taylor series expansion of an expected utility index demonstrates that risk is to be measured in terms of reward variance; all other aspects of risk (skewness, kurtosis, etc.) are secondary. See (Bali, Demirtas et al. 2006).

⁴ A few neuroscientists have started to incorporate risk in reinforcement learning. They would refer to the (prediction) risk as *expected uncertainty* Yu, A. J. and P. Dayan (2003). "Expected and unexpected uncertainty: ACh and NE in the neocortex." Advances in Neural Information Processing Systems **15**.. For a rigorous treatment, see [Preuschoff-Bossaerts].

The picture that emerges from the extant findings is one where the brain values risky gambles by evaluating their expected reward and risk separately (Figure 5). The separate evaluations are then merged to generate a total valuation signal, detectable in, e.g., prefrontal cortex (PFC). Figure 6, adapted from (Tobler, O'Doherty et al. 2007), shows how PFC activation increases with expected reward for all subjects, but decreases with risk for risk-averse subjects and increases with risk for risk-seeking subjects.

Perhaps unexpectedly, the separation between expected reward and variance may also carry over to choices that require tradeoffs between efficiency and inequity. The latter is a typical problem in studies of distributive justice. A deep theoretical connection exists between decision-making under uncertainty and the measurement of inequity (Atkinson 1970). The central concern in both is the comparison of frequency distributions $f(y)$. In decisions under uncertainty, $f(y)$ denotes the probability of state y occurring, whereas in the case of inequity, $f(y)$ is the weight placed on the income or other variables of interest for person y .

In experiments involving efficiency and inequity tradeoffs in a distributive justice task, Hsu et al. (2007) found a separation of regions encoding efficiency and inequity in much the same way as that of expected reward and variance. Specifically, efficiency is encoded in striatal regions including the caudate and putamen, whereas inequity is encoded in the insular cortex. This utilization of similar neural mechanisms for what on the surface are very different decisions provide a strong indication that these are very general mechanisms for the computation of decisions of a variety of different situations.

Just like encoding of expected reward presumes (expected) reward learning, the fact that the brain tracks risk suggests that it somehow learns about it. To date, very little is known about this. Still, (Preuschoff, Quartz et al. 2007) recently discovered that phasic activation in insula correlates with risk prediction errors, i.e., with the difference between realized risk (the size of the reward prediction error) and anticipated risk (the expected size of the reward prediction error). In their paradigm, risk prediction errors occur at two points in a trial; activations appear to be consistent at both points in time.

One may object that one can record a risk prediction error in their paradigm. Indeed, the risk prediction error in a trial is irrelevant to make better forecasts in future trials, as the trials are independent. That is, there is no relevant learning. There are three reactions to this objection. First, this finding is not unique to risk prediction errors. One can similarly record reward prediction errors in situations where learning is irrelevant because trials are independent (see, e.g., (Fiorillo, Tobler et al. 2003), (Preuschoff, Bossaerts et al. 2006)). Second, encoding of prediction errors is relevant for cases where learning unexpectedly does matter. If the experimenter told the subject that trials are independent [as in (Preuschoff, Bossaerts et al. 2006)] yet the experimenter lied [unlike in (Preuschoff, Bossaerts et al. 2006)], the brain better have prediction errors at hand to better manage the situation⁵. In this sense, the emergence of prediction errors in situations where learning is irrelevant in principle is equivalent to the presence of valuation signals in imperative trials.

⁵ In fact, this is one of the prime reasons why deception is frowned upon in economic experiments, but it would not be a compelling one if people always blindly believed the experiment.

Third, the prediction error encoding sometimes does have behavioral implications. Choices across independent trials sometimes reflect belief updating as if outcomes are actually dependent. Such belief updating presumes encoding of prediction errors. A classical example is underweighting of small probability events in experiments with trial-by-trial reward feedback [see (Hertwig, Barron et al. 2004)].

Finally, risk encoding may play a role in learning in addition to guiding choice. For learning to be effective, however, risk encoding should be *objective*. In particular, risk aversion or risk tolerance should not affect risk perception. Imagine that more risk-averse agents over-estimate risk. Since perceived risk decreases the learning rate, this would mean that risk-averse agents learn more slowly. That would give them a handicap in competition with agents who are more risk tolerant. To date, it is not known to what extent learning is affected by risk aversion, or to what extent the brain encodes a dual risk signal – one subjective, to guide choice, and one objective, to drive learning.

4. Extending The Reward-Risk Computational Model To Ambiguity

So far we have considered only pure risk, i.e., situations where probabilities are known, perhaps after a lengthy learning episode, or because probabilities are just given. The evidence points towards a computational model whereby the brain computes value by separately encoding expected reward and risk, and combining the results. Such a computational model is known to approximate well many types of utility functions (Bali, Demirtas et al. 2006) including prospect theory (Agren 2006).

In many situations in real-life, however, probabilities are unknown or partially known. This is called ambiguity in decision theory (Ellsberg 1961). A stylized fact in decision experiments is that many individuals are ambiguity averse (Camerer and Weber 1992). A number of models exist that account for decisions under ambiguity. Perhaps the best known is the maxmin utility model of Gilboa and Schmeidler (Gilboa and Schmeidler 1989), and its extension the α -maxmin utility model (Ghirardato, Maccheroni et al. 2004); (Huettel, Stowe et al. 2005). In the latter, the worst-case and best-case scenarios provide anchor points for beliefs. In one extreme case ($\alpha=1$), beliefs are entirely determined by the worst-case scenario, and the α -maxmin utility model coincides with that of (Gilboa and Schmeidler 1989). In the other extreme case ($\alpha=0$), beliefs are solely based on the best-case scenario. For alpha larger than 0.5, the decision maker is ambiguity averse; for alpha less than 0.5, she is ambiguity loving; at alpha equal to 0.5, she is ambiguity neutral; her beliefs are the average of the worst-case and best-case scenarios.

For instance, the decision maker may bet on the drawing of red, green and blue balls from an urn. She earns \$1 if the ball drawn is red or green; if the ball is blue, she loses \$1. One third of the balls are red. The number of green and blue balls is unknown, but the decision maker knows that at least 1/6 of the balls are green, and 1/12 are blue.

A decision maker with $\alpha=0$ will only consider the worst-case scenario: the minimum probability of winning p_{min} equals $\frac{1}{2}$ ($=1/3+1/6$); this will fix her beliefs. When $\alpha=1$, the maximum probability p_{max}

determines her beliefs; it equals $11/12 (=1-1/12)$. In general, the belief that the decision maker wins equals $\alpha p_{min} + (1-\alpha) p_{max}$. Without loss of generality, we can set the utility of winning \$1 equal to 1 and the utility of losing \$1 equal to 0. Hence, the general expression of the expected utility of the decision maker is $\alpha p_{min} + (1-\alpha) p_{max}$.

As with expected utility under pure risk, the α -maxmin utility model is equivalent (this time, it is *not an approximation*) to one where utility is computed based on a trade-off between mean and risk. In this case, the mean is just the average of the minimum and maximum probabilities:

$$\bar{p} = \frac{1}{2}p_{min} + \frac{1}{2}p_{max}.$$

The risk c , in contrast, is determined by the amount ambiguity:

$$c = p_{max} - p_{min}.$$

Utility, then, is some weighted average of \bar{p} and c :

$$U = \gamma_1 \bar{p} + \gamma_2 c.$$

Simple algebra reveals that:

$$\alpha = \frac{1}{2} \gamma_1 - \gamma_2,$$

$$1 - \alpha = \frac{1}{2} \gamma_1 + \gamma_2.$$

So $\gamma_1 = 1$ and $\gamma_2 = 1/2 - \alpha$.

Therefore, under both pure risk and ambiguity, there exists an equivalent representation of utilities in terms of a model that trades off mean against risk. As such, one can conjecture that brain regions involved in encoding mean and risk in situations of pure risk (striatum, anterior cingulate cortex, insula, inferior frontal gyrus,...) are also involved in encoding these parameters when there is ambiguity. As the weights on mean and risk need not be the same, it is to be expected that the intensity of the brain activations differs depending on the situation. This is consistent with the findings of (Hsu, Bhatt et al. 2005) and (Huettel, Stowe et al. 2006).

Closer inspection reveals that the mean-risk model is actually more general than the α -maxmin model: the latter imposes restrictions on the weights γ_1 and γ_2 in the mean-risk model. In fact, the mean-risk mode could be derived as a hierarchical Bayesian updating model, whereby the decision maker is first agnostic about the true probability of winning (in the example above, it is uniform between $1/3$ and 1), but then uses the midpoint (average between minimum and maximum probabilities) as a cue where the true probability lies. Specifically, the midpoint is assumed to be drawn from a uniform distribution centered at the true probability and with range equal to the difference between the maximum and

minimum probabilities. The posterior mean and variance of the true probability are given by \bar{p} and c , respectively.

The connection with updating and learning is important and illuminating in light of findings of amygdala activation in (Hsu, Bhatt et al. 2005). Significant activation of amygdala is rare or non-existent under pure risk. When contrasting choices involving ambiguity against those involving only pure risk, however, amygdala activation is strong. This is consistent with the amygdala's role in learning, for example, in oddball and novelty tasks (Rutishauser, Mamelak et al. 2006), and suggests that the amygdala signals a need for learning of the unknown probabilities. Further research is needed to discover the true role of amygdala under ambiguity.

More generally, further experiments are needed to verify to what extent valuation in the context of ambiguity is based on a computational model that trades off mean against risk. These experiments should extend previous studies pairing ambiguous with pure-risk (or zero-risk) gambles Hsu [(Hsu, Bhatt et al. 2005), (Huettel, Stowe et al. 2006)], and focus on choice between gambles with differing levels of ambiguity.

Figure 1

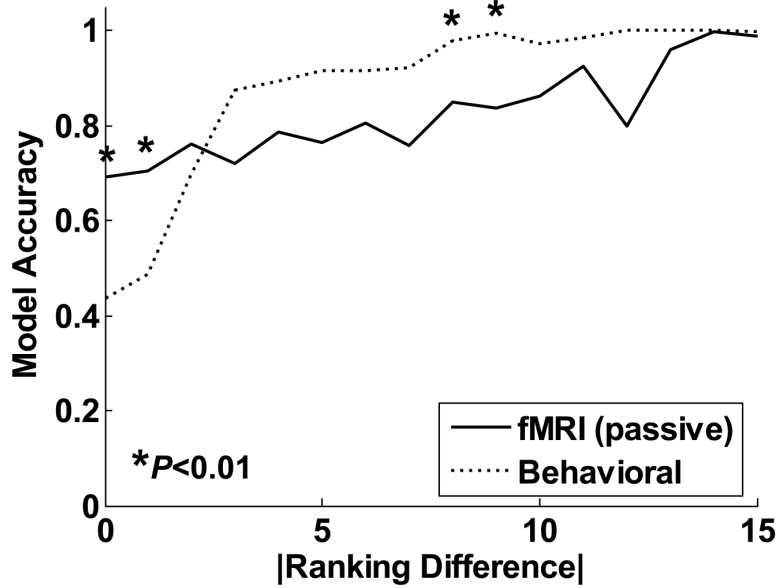


Figure 1 Caption: Out-of-sample prediction accuracy (percentage correct) of two models of binary choice, as a function of the ranking difference of the two options revealed through choice. The fMRI (passive) model uses as inputs activation in brain regions identified to be involved in valuation during purely imperative trials. The Behavioral model is based on rankings revealed in choice in free-choice trials. The parameters of both models are fit on free choice trials. Prediction accuracy is out-of-sample, i.e., it is measured in free choice trials not used for fitting the models. Neural activation identified in imperative trials (the fMRI model) predicts choices better than actual choice in free-choice trials when the two available options are close in ranking. *Source:* (Berns, Capra et al. 2007).

Figure 2.

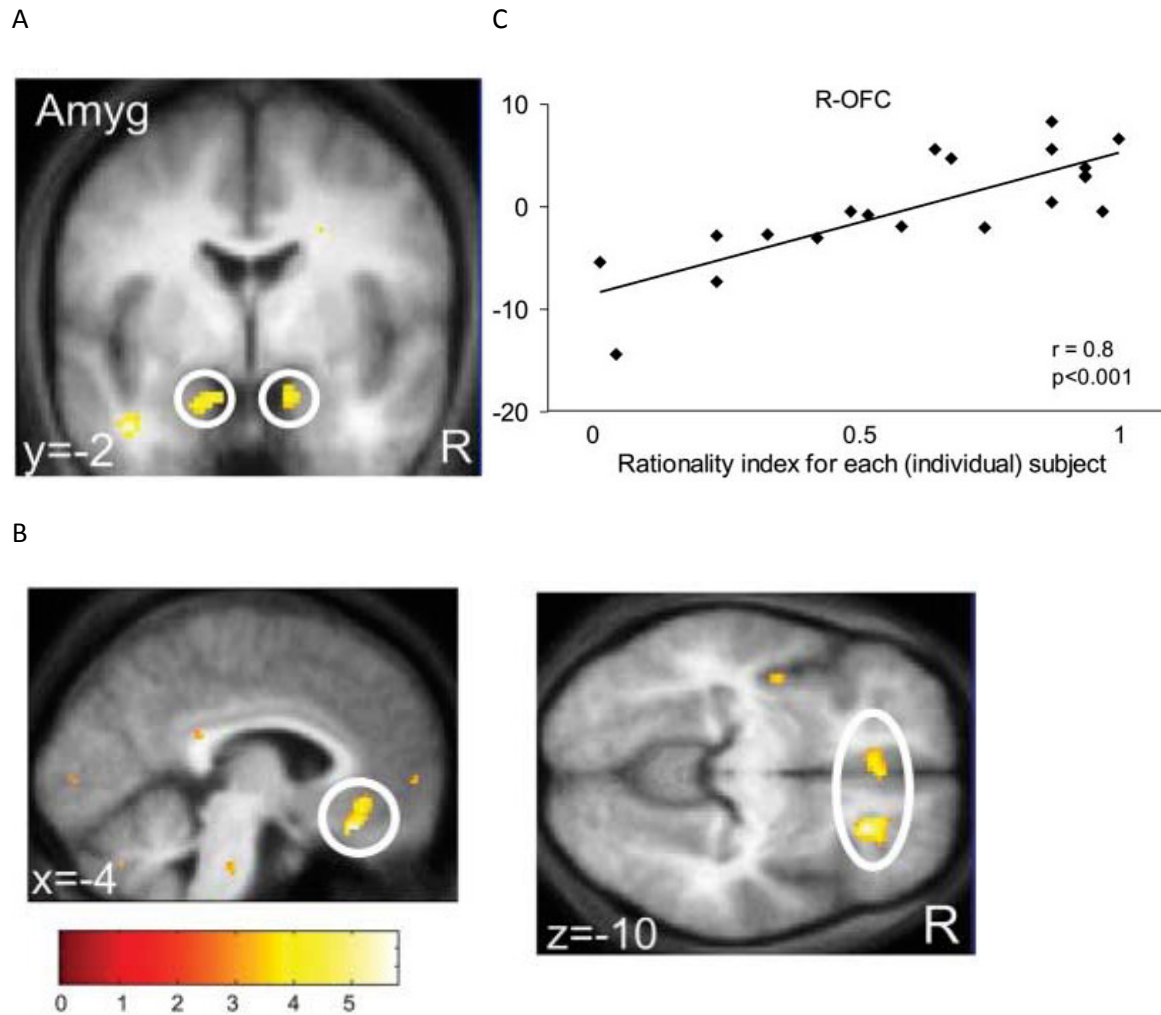


Figure 2 Caption: A: Amygdala activation correlates with contrast between lotteries cast in terms of gains and in terms of losses. B, C: Contrast in prefrontal cortex activation between lotteries cast in terms of gains and in terms of losses increases with subjects' rationality index. This index is measured as the difference between the proportion of trials in which subjects chose the gamble in the loss frame, as compared to the gain frame, and then linearly transformed such that 1 is the most rational. *Source:* (De Martino, Kumaran et al. 2006).

Figure 3

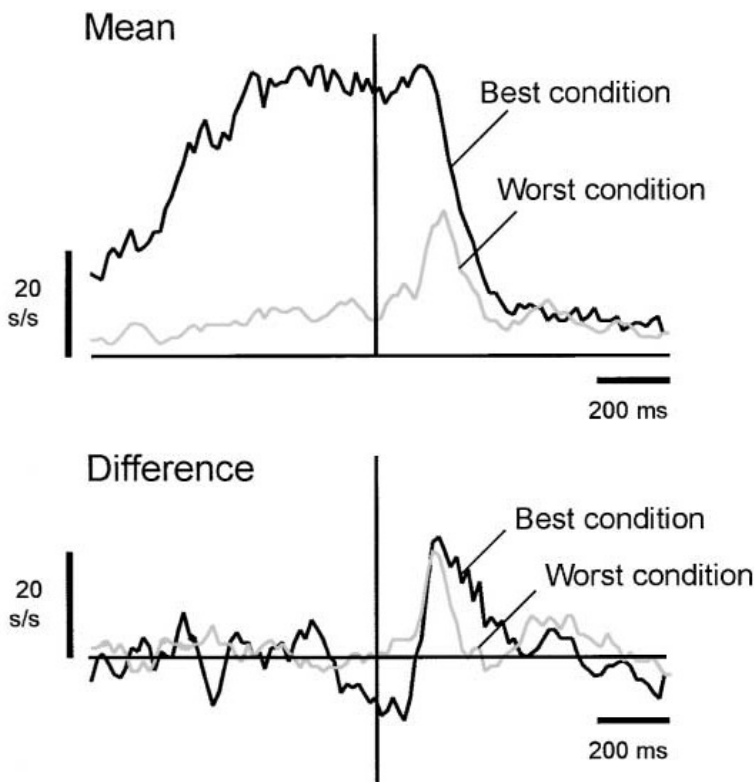


Figure 3 Caption: Average firing rates of a cluster of caudate neurons in the monkey brain that encode the value of eye movements (saccades) in the preferred direction. In trials where this direction generates higher reward on average (“Best condition”), increase in firing starts several hundred milliseconds (ms) before stimulus emerges (vertical line) that indicates which direction will be rewarded (top graph, labeled “mean”). Anticipatory increase in firing rates are absent in trials where the preferred direction receives the lower reward on average (“Worst condition”). After stimulus presentation, an increase is recorded in average firing rates across trials where the stimulus indicates reward for a saccade in the preferred direction, relative to average firing rates when saccades are instructed in the opposite direction. The difference (bottom graph) is the same for both “Best condition” and “Worst condition” trials. *Source:* (Lauwereyns, Takikawa et al. 2002).

Figure 4

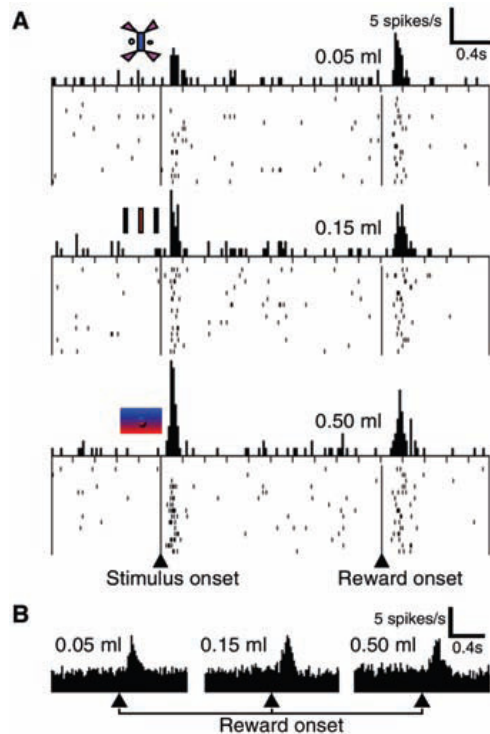


Figure 4 Caption: Single-trial firing rates and corresponding histogram for a single dopamine neuron in ventral tegmental area (A) and histogram of group firing rates (B) in imperative trials where a monkey is presented with stimuli that indicate the size of a random (50% chance) juice reward; results are shown only for rewarded trials. Firing at stimulus onset reflects the size of the expected juice reward. Firing at reward onset reflects the prediction error (positive, since only rewarded trials are retained for the analysis). The firing at reward delivery does not, however, reflect the size of the prediction error, which increases with the size of the reward (e.g., 0.50ml against an expectation of 0.25ml in the third case, or 0.05ml against an expectation of 0.025ml in the first case). Firing at reward delivery reflects a scaled prediction error: the prediction error divided by the anticipated size of the prediction error (prediction risk). *Source:* (Tobler, Fiorillo et al. 2005).

Figure 5

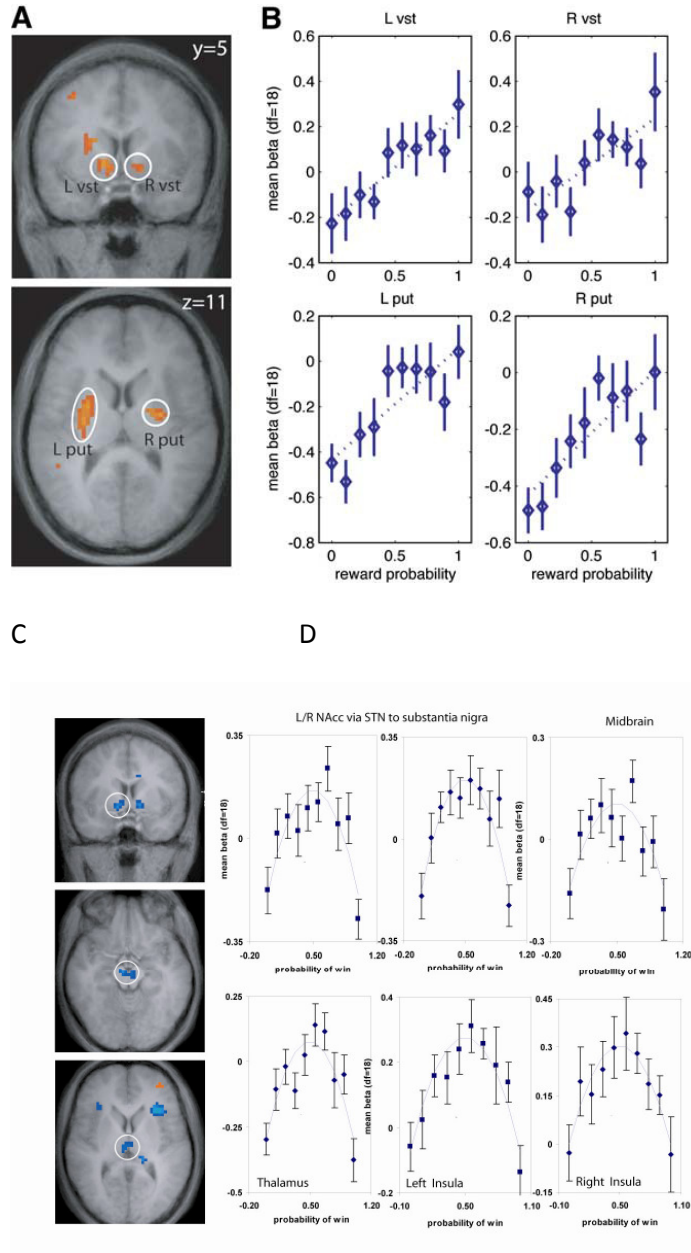


Figure 5 Caption: (Immediate) activation in striatal areas (A) increases with the probability of a fixed reward (B) in the human brain in gambles where the probability of reward fluctuates between 0 and 1. As such, activation correlates with expected reward. (Delayed) activations in sub-cortical areas and insula (C) change quadratically with probability of reward (D), with a peak at 50% chance of winning, and minima at certainty of loss or gain. As such, activation correlates with risk, measured as reward variance.

Vertical line segments indicate 95% confidence intervals. *Source:* (Preuschoff, Bossaerts et al. 2006), and unpublished data.

Figure 6

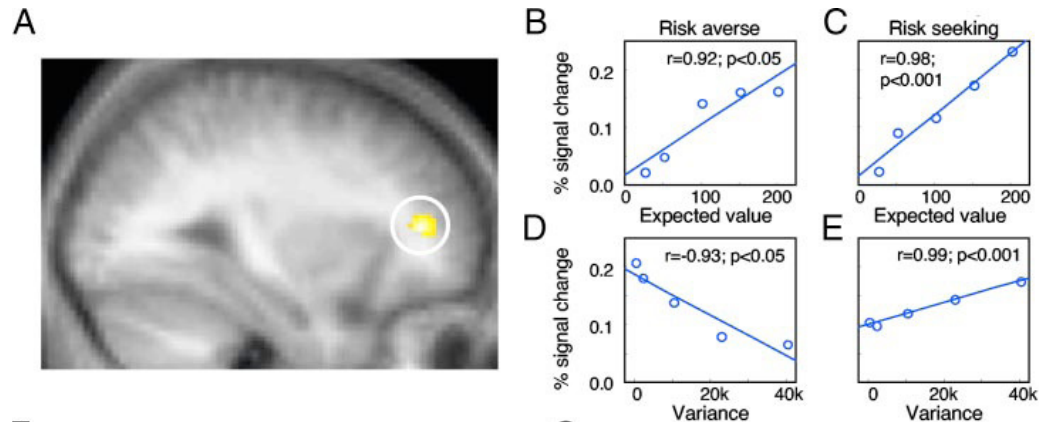


Figure 6 Caption: Activation in medial prefrontal cortex correlates with expected utility of a gamble in imperative trials (A). This activation increases in expected reward both for risk averse (B) and risk seeking (C) subjects; risk attitudes are measured through revealed preference in free-choice trials. The same activation decreases with risk of the gamble (measured as reward variance) for risk averse subjects (D). It increases with risk for risk seeking subjects (E). *Source:* (Tobler, O'Doherty et al. 2007).

Figure 7

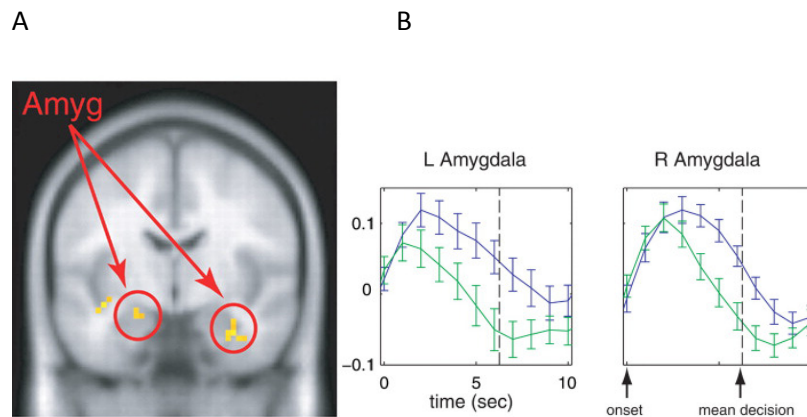


Figure 7 Caption: Activation in amygdala is higher in free-choice trials involving an ambiguous option than those only involving pure risk. (A) Location of activation. (B) fMRI activation as a function of time since trial onset; dashed vertical line indicates mean time of choice; blue line is average activation in trials with ambiguity; green line is average activation in trials with pure risk; vertical line segments indicate 95% confidence intervals. Autocorrelation in time series is meaningless, as is scale of activation; only the difference in activation matters. *Source:* (Hsu, Bhatt et al. 2005).

References

- Agren, M. (2006). "Prospect Theory and Higher Moments." Working Paper, Uppsala University.
- Atkinson, A. B. (1970). "On the measurement of inequality." Journal of Economic Theory **2**(3): 244-263.
- Bali, T. G., K. O. Demirtas, et al. (2006). "Is There a Relation Between Downside Risk and Expected Stock Returns?" Working Paper.
- Bechara, A. and A. R. Damasio (2005). "The somatic marker hypothesis: A neural theory of economic decision." Games and Economic Behavior **52**: 336-372.
- Bechara, A., H. Damasio, et al. (1999). "Different Contributions of the Human Amygdala and Ventromedial Prefrontal Cortex to Decision-Making." J. Neurosci. **19**(13): 5473-5481.
- Bechara, A., H. Damasio, et al. (1997). "Deciding Advantageously Before Knowing the Advantageous Strategy." Science **275**(5304): 1293-1295.
- Bechara, A., H. Damasio, et al. (1997). "Deciding Advantageously Before Knowing the Advantageous Strategy." **275**: 1293-1295.
- Behrens, T. E., M. W. Woolrich, et al. (2007). "Learning the value of information in an uncertain world." Nat Neurosci.
- Bernheim, B. D. and A. Rangel (2006). "Beyond Revealed Preference: Toward Choice-Theoretic Foundations for Behavioral Welfare Economics." Working Paper.
- Berns, G. S., M. Capra, et al. (2007). "Nonlinear neurobiological probability weighting functions for aversive outcomes." NeuroImage(Forthcoming).
- Blume, L. and D. Easley (1992). "Evolution and market behavior." Journal of Economic Theory **107**: 95-130.
- Camerer, C. and M. Weber (1992). "Recent Developments in Modeling Preferences - Uncertainty and Ambiguity." Journal of Risk and Uncertainty **5**(4): 325-370.
- Critchley, H., C. Mathias, et al. (2001). "Neural activity in the human brain relating to uncertainty and arousal during anticipation." Neuron **29**(2): 537.
- De Martino, B., D. Kumaran, et al. (2006). "Frames, Biases, and Rational Decision-Making in the Human Brain." Science **313**(5787): 684-687.
- Diaconis, P. and D. Freedman (1986). "On the consistency of Bayes estimates." The Annals of Statistics **14**(1): 1-26.
- Ellsberg, D. (1961). "Risk, Ambiguity, and the Savage Axioms." Quarterly Journal of Economics **75**(4): 643-669.
- Fiorillo, C. D., P. N. Tobler, et al. (2003). "Discrete coding of reward probability and uncertainty by dopamine neurons." Science **299**(5614): 1898-902.
- Ghirardato, P., F. Maccheroni, et al. (2004). "Differentiating Ambiguity and Ambiguity Attitude." Journal of Economic Theory **118**: 133-173.
- Gilboa, I. and D. Schmeidler (1989). "Maxmin expected utility with non-unique prior." Journal of Mathematical Economics **18**(2): 141-153.
- Gomes, F. J. (2005). "Portfolio Choice and Trading Volume with Loss-Averse Investors." Journal of Business **78**: 675-706.
- Hakansson, N. H. (1971). "Multi-period mean-variance analysis: Toward a general theory of portfolio choice." Journal of Finance **26**: 857-884.

- Hertwig, R., G. Barron, et al. (2004). "Decisions from experience and the effect of rare events in risky choice." Psychological Science **15**(8): 534-539.
- Hollerman, J. R. and W. Schultz (1998). "Dopamine neurons report an error in the temporal prediction of reward during learning." Nat Neurosci **1**(4): 304-9.
- Hsu, M., C. Anen, et al. (2007). The Right and the Good: Distributive Justice and Neural Encoding of Equity and Efficiency. Working Paper.
- Hsu, M., M. Bhatt, et al. (2005). "Neural systems responding to degrees of uncertainty in human decision making." Science **310**(5754): 1680-1683.
- Huettel, S., A. Song, et al. (2005). "Decisions under Uncertainty: Probabilistic Context Influences Activation of Prefrontal and Parietal Cortices." J. Neurosci. **25**(13): 3304-3311.
- Huettel, S., C. Stowe, et al. (2005). "Choices between gambles: Effects of certainty, risk, and ambiguity upon brain systems for decision making and reward evaluation." Journal of Cognitive Neuroscience: 221-221.
- Huettel, S. A., C. J. Stowe, et al. (2006). "Neural Signatures of Economic Preferences for Risk and Ambiguity." Neuron **49**: 765-775.
- Kahneman, D. and A. Tversky (1992). "Advances in Prospect Theory: Cumulative Representation of Uncertainty." Journal of Risk and Uncertainty **5**: 297-324.
- Lauwereyns, J., Y. Takikawa, et al. (2002). "Feature-Based Anticipation of Cues that Predict Reward in Monkey Caudate Nucleus." Neuron **33**: 463-473.
- Lucas, R. (1976). "Econometric Policy Evaluation: A Critique." Carnegie-Rochester Conference Series on Public Policy **1**: 19-46.
- Marinacci, M. (2002). "Probabilistic sophistication and multiple priors." Econometrica **70**: 755-764.
- McClure, S. M., G. S. Berns, et al. (2003). "Temporal prediction errors in a passive learning task activate human striatum." Neuron **38**(2): 339-46.
- McFadden, D. (1974). Conditional Logit Analysis of Qualitative Choice Behavior. Frontiers in Econometrics. P. Zarembka. New York, Academic Press: 105-142.
- Montague, P. R., P. Dayan, et al. (1996). "A Framework for Mesencephalic Dopamine Systems Based on Predictive Hebbian Learning." Journal of Neuroscience **16**: 1936-1947.
- O'Doherty, J. P., P. Dayan, et al. (2003). "Temporal difference models and reward-related learning in the human brain." Neuron **38**(2): 329-37.
- Paulus, M. P., C. Rogalsky, et al. (2003). "Increased activation in the right insula during risk-taking decision making is related to harm avoidance and neuroticism." NeuroImage **19**: 1439-1448.
- Preuschoff, K. and P. Bossaerts (2007). "Adding prediction risk to the theory of reward learning." Ann N Y Acad Sci **1104**: 135-46.
- Preuschoff, K., P. Bossaerts, et al. (2006). "Neural differentiation of expected reward and risk in human subcortical structures." Neuron **51**(3): 381-90.
- Preuschoff, K., S. R. Quartz, et al. (2007). "Human Insula Activation Reflects Risk Prediction Errors As Well As Risk." Under Review.
- Rutishauser, U., A. N. Mamelak, et al. (2006). "Single-Trial Learning of Novel Stimuli by Individual Neurons of the Human Hippocampus-Amygdala Complex." Neuron **49**(6): 805-813.
- Segal, U. (1987). "The Ellsberg Paradox and Risk-Aversion - an Anticipated Utility Approach." International Economic Review **28**(1): 175-202.
- Tobler, P. N., C. D. Fiorillo, et al. (2005). "Adaptive coding of reward value by dopamine neurons." Science **307**(5715): 1642-5.
- Tobler, P. N., J. P. O'Doherty, et al. (2007). "Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems." J Neurophysiol **97**(2): 1621-32.
- Yu, A. J. and P. Dayan (2003). "Expected and unexpected uncertainty: ACh and NE in the neocortex." Advances in Neural Information Processing Systems **15**.

