

Neurobiological studies of risk assessment: A comparison of expected utility and mean–variance approaches

MATHIEU D'ACREMONT

Swiss Federal Institute of Technology, Lausanne, Switzerland

AND

PETER BOSSAERTS

Swiss Federal Institute of Technology, Lausanne, Switzerland
and California Institute of Technology, Pasadena, California

When modeling valuation under uncertainty, economists generally prefer expected utility because it has an axiomatic foundation, meaning that the resulting choices will satisfy a number of rationality requirements. In expected utility theory, values are computed by multiplying probabilities of each possible *state of nature* by the payoff in that state and summing the results. The drawback of this approach is that all state probabilities need to be dealt with separately, which becomes extremely cumbersome when it comes to learning. Finance academics and professionals, however, prefer to value risky prospects in terms of a trade-off between expected reward and risk, where the latter is usually measured in terms of reward variance. This mean–variance approach is fast and simple and greatly facilitates learning, but it impedes assigning values to new gambles on the basis of those of known ones. To date, it is unclear whether the human brain computes values in accordance with expected utility theory or with mean–variance analysis. In this article, we discuss the theoretical and empirical arguments that favor one or the other theory. We also propose a new experimental paradigm that could determine whether the human brain follows the expected utility or the mean–variance approach. Behavioral results of implementation of the paradigm are discussed.

When choices under uncertainty satisfy certain basic rationality criteria, they can be thought of as maximizing a utility index that is obtained by multiplying probabilities of possible states by utilities of the outcomes promised in each of the states. This equivalence between choices and maximization of an *expected utility index* was first demonstrated by Von Neumann and Morgenstern (1947) and has since been proven to hold under quite general conditions of uncertainty (Savage, 1972) and far less stringent conditions of rationality—better reflecting the properties of actual human choices (Tversky & Kahneman, 1992).

The maximization of an expected utility index may be thought of not only as representing choices, but also as a *means to compute choice*. This view is implicit in much of recent neuroeconomic work in which attempts are made to find separate neurocorrelates of probabilities (Chandrasekhar, Capra, Moore, Noussair, & Berns, 2008), of utilities assigned to magnitudes (Tom, Fox, Trepel, & Poldrack, 2007), or both (Tobler, O'Doherty, Dolan, & Schultz, 2007). But although expected utility theory does provide an effective way to compute choices under uncertainty, it is by no means the only way.

Indeed, in financial economics, it has long been tradition to compute the value of risky gambles in terms of

statistical moments, expected payoff, payoff variance, and so forth (Black & Scholes, 1973; Markowitz, 1952).

The approach is not unrelated to expected utility: A mathematical operation called the *Taylor series expansion* demonstrates that a finite number of moments suffices to approximate well any smooth expected utility index (see the Appendix). This being said, we hasten to add that financial economists usually consider only the first two statistical moments (namely, expected payoff and payoff variance); two moments would, in general, provide only a very crude approximation of expected utility. Also, the square root of variance—that is, standard deviation—is often used as a measure of risk, instead of variance, because in many realistic cases (examples will follow), standard deviation is of the same order of magnitude as mean payoff and, hence, easily comparable.

Here, we raise the fundamental question of how the human brain computes choices. Does it follow the approach in classical decision theory, multiplying state probabilities by utilities of magnitudes to be received in each state, or does it rather opt for the financial approach, assessing expected reward and risk (measured as variance), to be integrated in a valuation signal that drives choices? The evidence for one or the other approach is far from

M. d'Acremont, mathieu.dacremont@epfl.ch

clear-cut. We mentioned a number of references that implicitly assume that the human brain separately encodes probabilities and (utilities of) magnitudes; there exist also studies that show separate encoding of expected reward and risk in the human brain (Preuschoff, Bossaerts, & Quartz, 2006; Tobler et al., 2007) and the nonhuman primate brain (Tobler, Fiorillo, & Schultz, 2005).

From a normative point, one may argue that the expected utility approach dominates, because it ensures that certain basic mistakes will never occur. Indeed, as we shall show, choices computed as a trade-off between mean (expected reward) and variance (of reward) may violate very basic principles of rationality. The violations are caused, however, not by the use of statistical moments to determine choice per se, but by the use of only a very limited number of moments (mean, variance). Efforts have also been made to avoid irrational choices by using other measures of risk, such as, for instance, the semivariance (variance computed with observations above or below the mean only; Ogryczak & Ruszczyński, 1999).

Nevertheless, computation of choice on the basis of the expected utility approach requires that the decision maker keep track of the probabilities of *all possible states*, a rather memory-intensive exercise. In contrast, the mean–variance approach requires encoding of only two basic numbers. Matters become even more complex when considering learning. In the absence of any specific link between the state probabilities (except that they add up to 1), the expected utility person has to update the probability of each state *separately*. If the number of states is large, the situation becomes untenable: Even with a (countably) infinite number of outcome observations, the expected utility decision maker will generically not learn the true probabilities (Diaconis & Freedman, 1986). Contrast this with a mean–variance decision maker who uses the simplest of learning algorithms—namely, the Rescorla–Wagner rule—to estimate the expected reward and risk (Rescorla & Wagner, 1972). This rule will generate precise estimates even after a few observations.

Expected utility theory and mean–variance preference theory differ also in the way they induce risk aversion. Expected utility theory predicts that risk aversion is related to the curvature of the utility function. The faster marginal utility decreases as a function of payoff, the higher risk aversion is. As such, one should be able to predict someone's risk attitudes from choices *not* involving any uncertainty. Specifically, to predict risk aversion, one merely has to observe a decrease in willingness to spend effort as the subject becomes richer. To date, nobody has convincingly been able to do so. The link between decreasing marginal utility and choice under uncertainty is absent in mean–variance choice theory. The link is also not as tight in prospect theory (Kahneman & Tversky, 1979), where weighing of the state probabilities interferes to induce risk aversion (see the Appendix).

The goal of this article is not only to clarify these issues, but also to propose a paradigm that could effectively determine whether choice is computed on the basis of the expected utility approach or on the basis of the mean–variance approach. The remainder is organized as follows.

We first will discuss how mean–variance–based choices may violate some basic principles of rationality. Subsequently, we will elaborate on the advantage of mean–variance representations of gambles for learning. We then will discuss recent empirical work that shows encoding of probability, magnitude, mean, and variance of reward in the human brain. We next will introduce a new paradigm to dissociate learning and decision making on the basis of expected utility and mean–variance analysis. We will discuss illustrative behavioral results based on implementation of the paradigm. We will close with some final remarks.

Mean–Variance Decisions May Lead to Dominated Choices

Decision theorists have long been interested in mathematical representations of choices that satisfy a minimum of rationality restrictions. These rationality restrictions are represented by means of *choice axioms*. Many expected utility decision-making models (including prospect theory) satisfy such choice axioms; Von Neumann–Morgenstern expected utility derives from the stringest axioms, whereas prospect theory is founded on far weaker axioms.

One of the weakest axioms is *statewise dominance*, also known as state-by-state dominance (see Harsanyi & Selten, 1988). Gamble A has strict statewise dominance over Gamble B if A gives a better outcome than does B in every possible state. The dominance is weak if A gives a better or equal outcome in every possible state, with strict inequality in at least one state. For instance, consider the following two Gambles A and B. Two states may occur, with a probability of .25 and .75, respectively (in fact, these probabilities are irrelevant for statewise dominance). Gamble A pays 5 in the first state and 3 in the second state. Gamble B pays 4 in the first state and 0 in the second state. So, Gamble A always pays more than Gamble B. As such, A statewise dominates B.

Statewise dominance is to be dissociated from stochastic dominance. The contrast between the two will, however, clarify the true meaning of expected utility theory. A Gamble C has (first-order) stochastic dominance over D if, for any outcome x , C gives a higher probability of receiving an outcome equal to or better than x under D. In other words, the cumulative distribution for Gamble C (F_C) is smaller than that for D (F_D)—that is, $F_C(x) < F_D(x)$. Statewise dominance is a stronger concept; in particular, it implies first-order stochastic dominance.

The fundamental difference is that statewise dominance requires the definition of states, which is not necessary for stochastic dominance. Thinking in terms of states is the foundation of probability theory and, by extension, expected utility theory. Under this perspective, randomness of a variable is a property of states; the variable is random only because it is a nontrivial function that maps states into payoffs. The concept of stochastic dominance does not require one to think in terms of states. Only (the distributions of) the outcomes are relevant. The same is true for mean–variance analysis.

When choosing an inferior, dominated gamble, the decision maker is said to make a *dominated choice*. The axiom of statewise dominance then says that rational agents

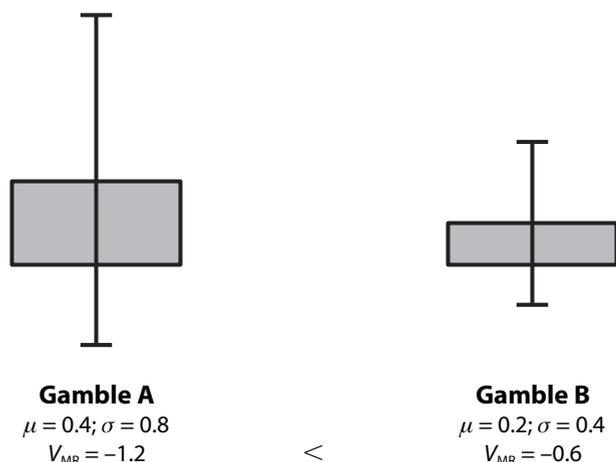


Figure 1. Representation of gambles within the mean-variance preference model. Gamble A dominates, but B is preferred.

should never make dominated choices. It is motivated by the premise that rational agents should always prefer more over less (*nonsatiation*). In terms of utility theory, the axiom states that utility must continue to increase with payoff. If the utility function u is differentiable, this means that the first derivative u' is positive everywhere.

The problem with choice based on a trade-off between expected (mean) payoff and variance (risk) of payoff is that it may violate the axiom of statewise dominance. Here is an example. Consider the following two-outcome gambles: A pays 2 and B pays 1 in the first state; both A and B pay zero in the second state; the first state has a probability of .20 of occurring; the second state occurs with a probability of .80. Obviously, A statewise dominates B. (To determine the statewise dominance, we did not need the probabilities, but we shall need those for the mean-variance analysis.)

How would a mean-variance optimizer evaluate these two gambles? The value V of a gamble for a mean-variance optimizer can be computed by subtracting the risk (measured as standard deviation σ) from the mean μ : $V_{MR} = \mu - b\sigma$. Here, the coefficient b measures the risk aversion; if b is high, our agent penalizes risk more and, hence, is more risk averse; if b is actually negative, our agent is risk seeking. Assume that $b = 2$. The mean payoff μ and standard deviation σ for A equal 0.4 and 0.8, respectively. For B, the numbers are 0.2 and 0.4. A graphical representation of the trade-offs between expected payoff and risk is given in Figure 1. Because $b = 2$, $V_{MR} = 0.4 - 2 \cdot 0.8 = -1.2$ for Gamble A and $V_{MR} = 0.2 - 2 \cdot 0.4 = -0.6$ for B. Thus, our agent assigns a higher value to B than to A and, hence, will decide to go for B. This choice is statewise dominated, however.

It can easily be shown that for any risk aversion parameter $b > 0.5$, B will be preferred to A. In other words, if a decision maker is sufficiently risk averse, he/she will prefer B because this gamble entails lower risk. In fact, it can be shown that statewise-dominated gambles can *always* be found for *any* mean-variance optimizer and that there are many of those (Borch, 1969).

By construction, expected utility theory does not violate the axiom of statewise dominance, provided the utility function is strictly increasing. Take someone with a simple utility function—namely, $u = \sqrt{x}$, where x denotes the payoff. The value of a gamble for this expected utility agent is $V_{EU} = E[u(x)] = E(\sqrt{x})$. The value of A is thus $V_{EU} = 0.20\sqrt{2} = 0.28$. The value of B is $V_{EU} = 0.20\sqrt{1} = 0.20$. The agent chooses A instead of the dominated lottery B. This can be shown to be true for any expected utility decision maker.

Figure 2 provides one way to graphically represent how an expected utility agent thinks about this particular problem. Gambles can be represented by circles. Circular segments denote states of the world; they are indexed by the symbol inside the segment (1 and 2). The size of a circular segment is proportional to the belief that the corresponding state will occur. The expected utility decision maker first compares the payoffs in each state. Payoffs are distinguished by color: Yellow denotes €0, blue €1, and orange €2. He or she will recognize that the payoff is the same in State 2, but higher for Gamble A in State 1. Hence, A is strictly preferred.

If humans are mean-variance optimizers, we should observe dominated choices. If expected utility is a better model of human decision making, dominated choices should be nonexistent. Numerous behavioral studies have reported that humans routinely violate this dominance principle (Birnbaum & Navarrete, 1998; Mellers, Berretty, & Birnbaum, 1995).

So, Why Mean-Variance Preferences?

The dominated choice argument is damaging for decision making based on mean-variance optimization. So, why would the brain choose to encode expected payoff and payoff variance, the building blocks of mean-variance preferences, instead of probabilities and utilities (magnitudes), the components needed to compute expected utility?

One argument in favor of representation of gambles in terms of mean payoff and payoff variance is that it facilitates learning. To see this, let us first go back to expected utility theory. This theory is based on the idea that there

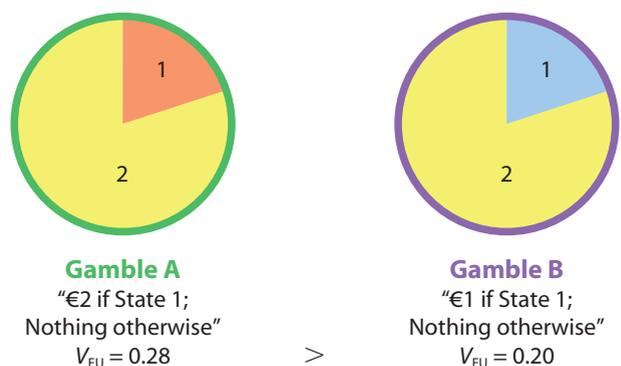


Figure 2. Representation of gambles within the expected utility model. States are identified by numbers (1 and 2); payoffs are identified by colors: yellow = €0, blue = €1, and orange = €2. Gamble A dominates and is preferred to Gamble B.

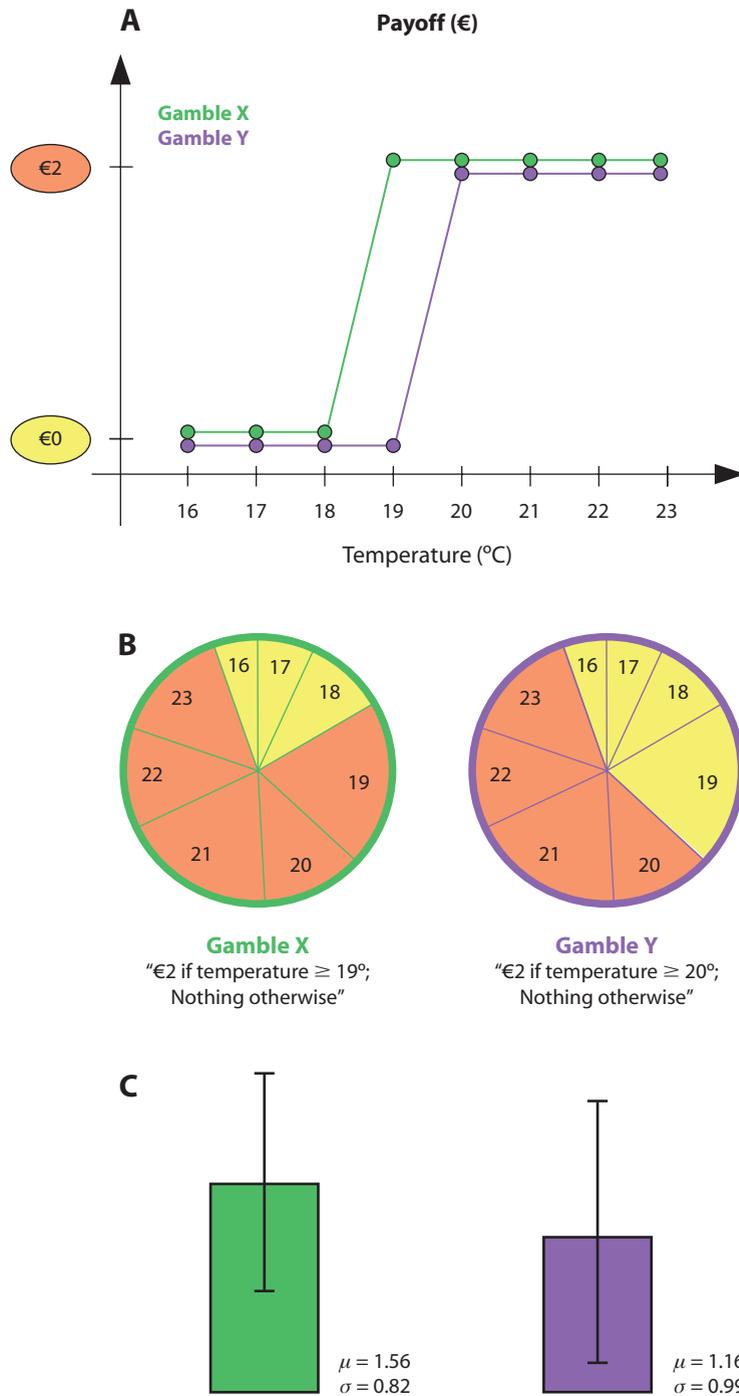


Figure 3. State-space conceptualization of random payoffs. (A) Maps between states and payoffs define gambles. (B) Expected utility representation of gambles. (C) For comparison: mean–variance representation of the same gambles.

are common, fundamental *states of nature* that underlie all gambles. The expected utility maximizer first has to ponder what these states of nature are and what probabilities to assign to each of them. When faced with a particular gamble, the expected utility maximizer will then determine how much this gamble pays in each of the possible

states. The value of the gamble is then computed by multiplying the utility of the payoff in each state by the belief that the state occurs and adding the results across states.

For instance, consider a Gamble X that pays €2 if the temperature tomorrow is equal or above 19°C and nothing otherwise and another Gamble Y that pays €2 if the tem-

perature tomorrow is equal to or above 20°C and nothing otherwise. Temperature levels can be considered to be the underlying states of nature for the gambles at hand, and the payoffs on the gambles can be represented as a mapping from these states to €0 or €2 (see Figure 3A). (For simplicity, we limit our attention to temperatures ranging from 16°C to 23°C.)

These two gambles can then be represented in a circle paradigm by letting the circular segments correspond to temperature levels; the temperature level inside a segment indicates what state the segment refers to. We now let color denote payoffs: yellow for €0 and orange for €2. Imagine now that our agent does not perfectly know the chances of each of the states (temperature levels). Our agent will learn as follows. He or she will record outcomes (temperatures) and use those to update his or her belief about the temperature levels. Effectively, each temperature level becomes a parameter, and the agent learns about each parameter separately.

One can easily see that such a learning strategy may become problematic. In our example, we limited our attention to eight temperature levels and, hence, seven parameters (one probability for each temperature level, except for the last one, where the probability can be inferred from the fact that probabilities should add up to 1). After, say, 7 days, our agent will have recorded only 1 observation on average per parameter, a rather dismal sample size. In realistic situations, the number of possible states and, hence, the number of parameters are obviously far higher. With 1,000 states, it takes 999 observations to obtain a ratio of observations to parameters equal to just 1. It becomes questionable whether this approach allows someone to learn anything at all! A formal treatment can be found in Diaconis and Freedman (1986).

As a result, the expected utility maximizer will have a hard time learning to correctly value risky gambles. What about the mean-variance optimizer? This agent is interested only in estimates of the mean payoff and of the payoff variance (or its square root, the standard deviation; see Figure 3C). These two parameters can easily be learned from observed outcomes, using simple reinforcement learning algorithms. The expected payoff can be estimated from the sample mean payoff of repeated play of the gamble, and the payoff variance can be estimated from the sample variance. With independent sampling, both the estimated mean and the variance will quickly converge to the true mean and variance, by the law of large numbers.

Note the fundamental difference in perspective between expected utility theory and mean-variance theory. In the former, all uncertainty about outcomes (payoffs) of different gambles can be reduced to the uncertainty of common states of nature. States of nature are effectively the deeper cause of outcome variation. In the latter, the decision maker considers only outcomes and learns about these.

There is an analogy between decision making based on the expected utility approach and learning based on Bayesian updating. Both perceive deeper causes behind observable outcomes. For the expected utility decision maker, these causes are referred to as states of nature. The Bayesian learner thinks of them as parameters.

Just as Bayesian learning can become a powerful inference tool in a complex world where outcomes may be generated by common causes (see, e.g., Hampton, Bossaerts, & O'Doherty, 2006), expected utility theory facilitates evaluation of uncertainty. Imagine, for instance, that the agent has always chosen Gamble X in Figure 3 in the past. The expected utility maximizer learns about the chances of winning by learning about the probabilities of each of the states (temperature levels). If asked what the chance is that Gamble Y will pay off, the expected utility maximizer will apply his or her knowledge of the probabilities of the temperature levels to induce an answer.

Compare this with the mean-variance optimizer. Since he or she has never seen any outcome for Gamble X, he or she has no idea about the expected payoff and payoff variance on Y. If he or she were to realize that there is a common cause (same states) behind the payoffs on Gambles X and Y, he or she would be able to obtain a good estimate of the expected payoff and the risk of Y on the basis of outcome observations for X.

The distinction between the expected utility and the mean-variance approaches to decision making is also analogous to the distinction between reflective and reflexive learning (Daw, Niv, & Dayan, 2005). The mean-variance approach is fast and easy, like reflexive learning, but subject to mistakes; the expected utility approach is slow and complicated, like reflective learning, but powerful. The mean-variance approach may be called *heuristic*, whereas the expected utility approach is formal (Kahneman & Frederick, 2002). The reflexive/reflective learning and heuristic/formal decision-making approaches are examples of the *dual-system* theory of cognition (Evans, 2003).

A final remark on the differences between the expected utility and the mean-variance theories concerns risk aversion. The two theories incorporate risk aversion (or risk seeking, for that matter) in fundamentally different ways. In expected utility, risk aversion is the consequence of nonlinearity in the valuation of magnitudes of outcomes. Specifically, when the utility function u of the payoff x is strictly concave in the payoff, the agent will exhibit risk aversion: He or she will favor risk-free gambles when alternative, risky gambles return the same payoff on average. In words, strict concavity means that the marginal utility of an extra payment declines with the payment amount. Mathematically, if u is twice differentiable in x , strict concavity obtains if the second derivative u'' is negative. The square root function [$u(x) = \sqrt{x}$] from the illustration in the previous section provides an example of a strictly concave function.

It may seem strange that risk aversion is determined solely by nonlinear valuation of outcomes. It is all the stranger because this nonlinearity has nothing to do with uncertainty. Specifically, risk aversion is *not* the result of biases in beliefs about the likelihood that states will occur.¹

In contrast, in mean-variance preference theory, risk aversion is the result of the penalty imposed on risk. Valuation V is based on a trade-off between mean μ and variance (standard deviation σ): $V_{MR} = \mu - b\sigma$, where b is the penalty imposed on risk. The higher b is, the more risk is

penalized and, hence, the higher the risk aversion exhibited in choices is.

One can interpret the multiplication $b\sigma$ in this valuation equation as a biasing of risk assessment σ by means of the coefficient b . As such, $b\sigma$ corresponds to *subjective risk assessment*. Hence, in mean–variance choice theory, risk aversion is the result of belief biases. In particular, risk aversion increases with overassessment of risk.

Simultaneous Encoding of Probability, Magnitude, Expectation, and Reward Variance

To date, it is not known whether the human brain computes value primarily on the basis of an expected utility approach or on the basis of a mean–variance approach. Neurobiological studies of brain activation in the context of risk generally analyze the data in terms of one or the other approach only. Studies that use the expected utility approach (Knutson, Taylor, Kaufman, Peterson, & Glover, 2005; Yacubian et al., 2006; Yacubian et al., 2007) correlate activation with probabilities and magnitudes (or utilities). For instance, Knutson et al. reported that effects of payoff magnitude and probability were preferentially observed in the nucleus accumbens and the medial prefrontal cortex, respectively. Studies that follow the mean–variance approach (Huettel, Song, & McCarthy, 2005; Paulus, Rogalsky, Simmons, Feinstein, & Stein, 2003; Preuschoff et al., 2006; Rolls, McCabe, & Redoute, 2008) identify neurocorrelates of expected reward and/or risk. For instance, expected value correlates with activity in the striatum, and risk correlates with activity in the striatum and insula (Preuschoff et al., 2006; Preuschoff, Quartz, & Bossaerts, 2008).

Studies that look at all aspects of decision making under uncertainty simultaneously (probability, magnitude, expected reward, and risk) are rare (for a review, see Knutson & Bossaerts, 2007). One important exception is Tobler et al. (2007). In this study, subjects observed abstract shapes representing gambles while brain activation was

recorded using fMRI. Levels of reward magnitude and probability were manipulated experimentally by changing number and shapes: The number of circles informed the subjects about reward magnitude, whereas color indicated probability (or the reverse, for the other half of the subjects). Effects of payoff magnitude, probability, and expected value were observed in distinct regions within the striatum and lateral prefrontal cortex (see Figure 4, reproduced from Tobler et al., 2007). A link with probability also emerged in the medial prefrontal cortex. Hemodynamic response in the lateral orbitofrontal cortex depended on an inverted-U function of probability, suggesting that this region may be involved in risk evaluation.

It is worth noting, however, that gambles have generally been represented with stimuli that separately convey probabilities and magnitudes. So, there is the possibility that the presence of activations correlating with probabilities and magnitudes is due to the way the information is presented from which the brain can infer expected reward and variance. That is, the brain may be “forced” to encode probabilities and magnitudes. Studies that analyze neuronal activity in terms of expected reward and variance have invariably presented gambles in the same way. Even when probabilities are indirectly conveyed—for example, through cards in a card game (Preuschoff et al., 2006)—the most immediate translation of the stimuli is to probabilities, whereas the mapping to payoff variance, in particular, is far from straightforward. As such, it remains unclear to what extent the human brain favors an approach to computing value of risky gambles that is based on expected utility theory or on mean–variance modeling.

As was stated before, in expected utility theory, risk sensitivity is caused by the curvature of the utility function. In mean–variance preference theory, risk aversion is the result of a penalty imposed on variance. This requires that variance be encoded separately from expected value. If the expected utility theory is true, neuroscientists should observe only brain activity related to the expected

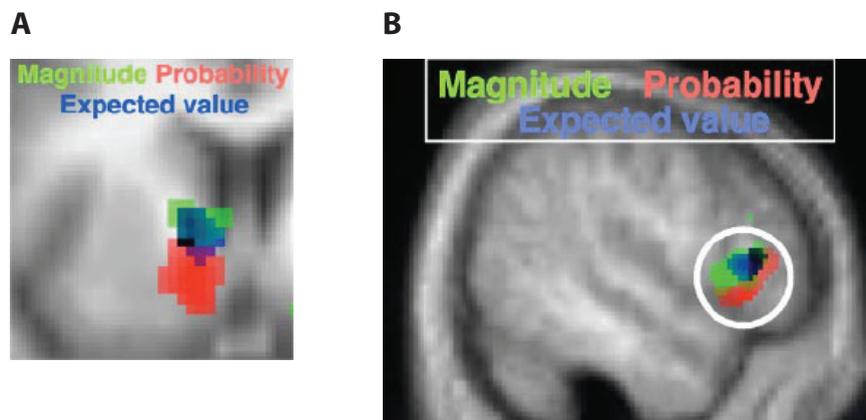


Figure 4. Brain activation related to payoff magnitude, probability, and expected value within the striatum (A) and the lateral prefrontal cortex (B). From “Reward Value Coding Distinct From Risk Attitude-Related Uncertainty Coding in Human Reward Systems,” by P. N. Tobler, J. P. O’Doherty, R. J. Dolan, and W. Schultz, 2007, *Journal of Neurophysiology*, 97, pp. 1626–1627. Copyright 2007 by The American Physiological Society. Reprinted with permission.

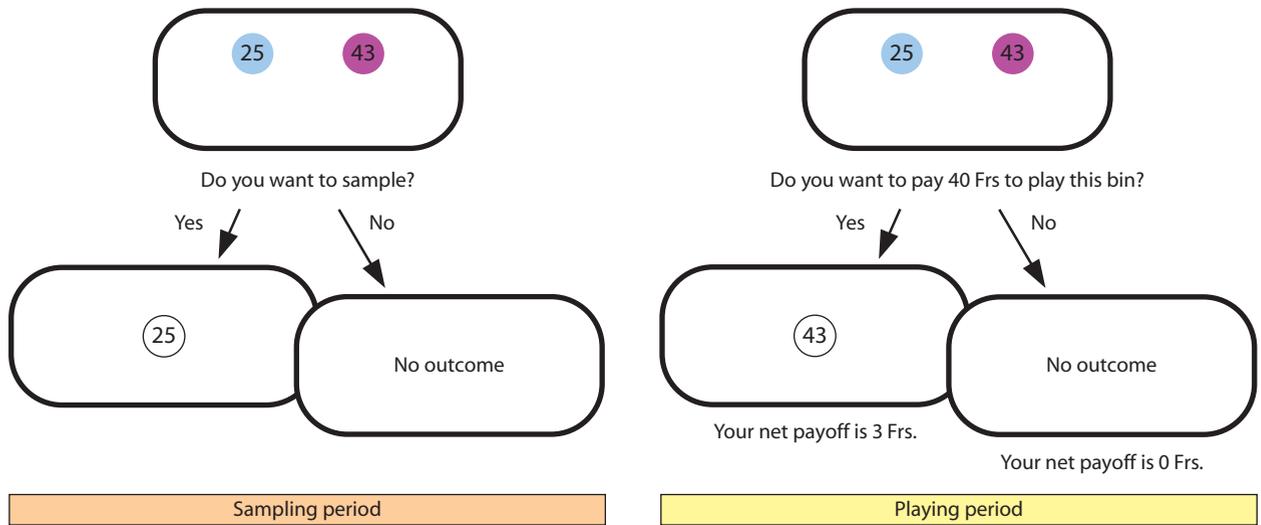


Figure 5. Sampling and playing periods within a trial.

utility of an option. If the second model is correct, two signals should emerge in the brain: one for expected value and the other for a risk measure such as variance. As it is, data slightly favor the latter hypothesis, because separate risk (variance) signals have been detected (Preuschoff et al., 2008; Tobler et al., 2007). Separate signals for expected value and variance have also been observed in a learning task. Indeed, a recent fMRI study has revealed that prediction error for expected value is related to activity in the striatum (a dopamine projection area generally understood to be involved in reward reinforcement learning; McClure, Berns, & Montague, 2003; O’Doherty et al., 2004), whereas prediction error for variance is related to activity in the insula and the inferior frontal gyrus (d’Acremont, Lu, Li, Van der Linden, & Bechara, 2007). Incidentally, the ubiquitous use of mean–variance analysis in the financial industry has led some to conjecture that this is actually the approach the human brain takes to determine choice under uncertainty (Dickhaut, 2008). The argument is that humans prefer to use computational tools that are based on principles employed by their brains.

A New Experimental Paradigm

Experimental design. To adjudicate between the expected utility and the mean–variance valuation principles, a new paradigm is needed that focuses on the two main differences—namely, potential of statewise dominated choice and, foremost, learning. We will propose one such paradigm here and will discuss the behavioral results of its implementation.

In our paradigm, states were represented by colors of balls in a bin, whereas payoffs (outcomes) were indicated by numbers on the balls. Payoffs were the same for all balls of the same color. Gambles were draws (with replacement) of a single ball from the bin; the number on the drawn ball determined the payoff. The subjects knew what colors (states) were in a bin at any moment and which payoff corresponded to each of the colors. However, the

subjects did not know the composition of the bin—that is, how many balls of each color there were.

We occasionally changed the bin. The period of time that we kept the bin the same was referred to as a *trial*. Each trial consisted of a *sampling period* and a *playing period* (see Figure 5).

Within a trial, the subjects were initially given the opportunity to sample the bin, as much as they liked, though at a small cost of 0.01 Frs (Swiss francs) per sample. With each sample, a ball was drawn from the bin (with replacement), and the payoff, but not the color of the ball, was revealed (see “Sampling period,” Figure 5).

Once the subjects decided not to sample anymore, we offered them a chance to buy into the gamble. We implemented a posted-price mechanism: We quoted a price, and the subject decided to take the offer or leave it. If the subject bought the gamble, we played it. We added the gamble payoff to, and subtracted the gamble price from, the play money. Otherwise, the subject kept the play money level. We repeated this five times with different prices, randomly chosen between the minimum and the maximum possible payoffs (see “Playing period,” Figure 5).

We preferred a posted-price mechanism for eliciting value over the usual incentive-compatible mechanisms whereby subjects are asked to offer a price. Our preference was motivated by ecological relevance. Indeed, most subjects rarely have to offer prices in daily life (when, say, buying milk in a supermarket); instead, prices are posted, and subjects decide whether to buy or not.

At the end of a trial, either we changed the composition (number of balls of each color and number of colors) and the payoffs on the balls (e.g., Trials 1 and 3 in Figure 6), or we changed *only* the payoffs (e.g., Trials 2 and 4 in the same figure). The subjects were told explicitly when this happened and which case applied.

Consequently, we effectively created a situation in which the subjects should have been fully aware of the underlying structure: They knew the states and the pay-

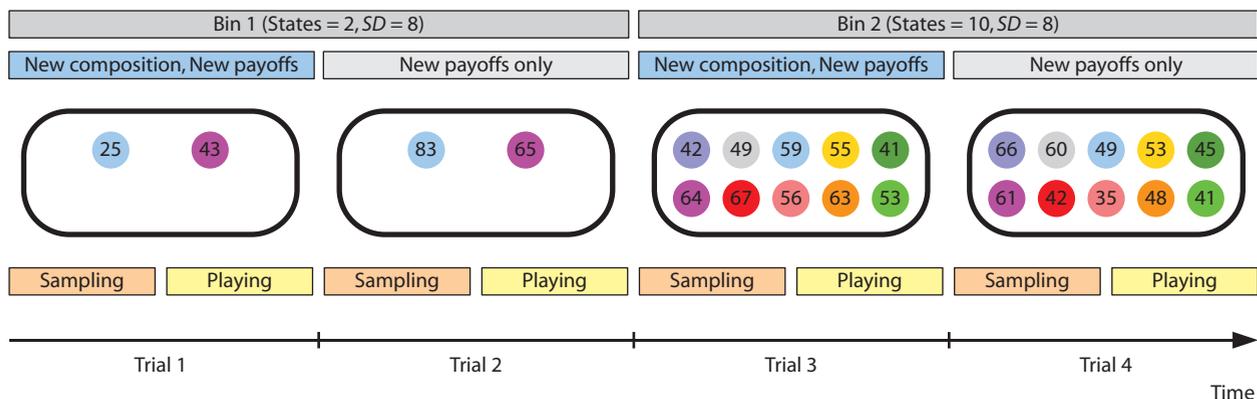


Figure 6. Time line of the experiment.

offs in each state. At the same time, all the subjects ever saw were payoffs; they never observed the states directly. They could have opted to ignore the states, effectively becoming mean–variance optimizers. Or they could have inverted payoffs for states, to apply the expected utility approach.

Bins were built in the following way. We first chose among the following number of states (colors): 2, 3, 5, or 10. A fixed set of probabilities (number of balls of a given color) was associated with each number of states. For instance, for three states, we chose $p_1 = .25$, $p_2 = .50$, and $p_3 = .25$. Subsequently, we assigned specific colors (from among 10 colors) to these probabilities. This defined a particular instance of a bin. Specific gambles were then constructed by assigning payoffs to colors in the bin in order to obtain a particular standard deviation (4, 8, or 12) and a mean chosen from a uniform distribution (ranging from 26 to 74 Frs).

Whenever we changed only the payoffs, we merely selected different payoffs for each of the colors, keeping the number of colors (states), the number of balls of each color (state probabilities), and the payoff standard deviation the same; only the mean payoff changed. Whenever we changed the composition of the bin, we chose a new number of colors (states) with its associated probabilities, assigned new colors to each of the probabilities, and picked a new payoff standard deviation.

Two different bins were built for each combination of number of states (colors) and payoff standard deviation. Thus, these two dimensions were orthogonalized. This makes in total 24 bins. For each of these bins, we changed payoffs once without changing the composition of the bin, to generate a different mean payoff. In total, we thus generated 48 trials per subject. The subjects were paid for performance. At the end of the experiment, they received one tenth of their net play money in real currency.

The interesting aspect of our paradigm is that expected utility and mean–variance models predict very different behavior. In particular, while sampling, expected utility subjects should infer from payoff which state has occurred, to keep track of the states (like Bayesians, they solve an inverse problem) and their frequency of occurrence. Therefore, two distinct features should characterize their decisions:

- E1. As long as the composition of the bin remains unchanged, expected utility maximizers will never resample when the mapping from states to payoffs (numbers on the balls) changes.
- E2. The length of sampling (number of times they sample) depends only on the number of states (number of colors in a bin).

Resampling is unnecessary (Feature E1) because expected utility decision makers know the composition of the bin (up to a certain level of precision). They need merely to apply the new mapping from states (colors) to payoffs in order to revalue gambles after a change in this mapping. The second feature presupposes that subjects stop sampling after their predictions reach a certain precision. Precision is measured in terms of the probabilities that a state (color) is drawn.

As was explained before, expected utility decision makers will never buy statewise dominated gambles. A simple implication for choice is the following:

- E3. Expected utility maximizers will never buy dominated gambles (where the price is at least as much as the highest possible payoff across all states) or pass on gambles dominated by their price (where the price is, at most, equal to the lowest possible payoff across all states).

In contrast, mean–variance optimizers focus on the distribution of payoffs, ignoring the underlying structure (mapping from states to payoffs). If they care only about the first two moments of the distribution, they will limit learning to the payoff mean and payoff variance. This they can accomplish using simple reinforcement learning algorithms such as Rescorla–Wagner (Rescorla & Wagner, 1972). But since they ignore the underlying states, mean–variance optimizers do need to relearn each time the mapping from states (colors) to payoffs changes, even if nothing about the states themselves (number of balls of each color) changes. Also, in determining how long to sample, the number of states is irrelevant for a mean–variance optimizer. Instead, the variance of the payoff is relevant (precision of estimates of the mean payoff depends on it), as is the kurtosis (the fourth moment, on which precision of

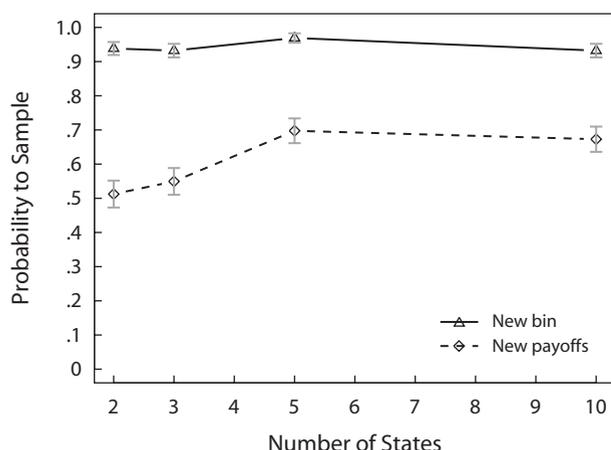


Figure 7. Mean probability of sampling as a function of number of states and whether the bin composition changes (“New bin”) or only the payoffs in each state (“New payoffs”). Vertical bars indicate standard errors.

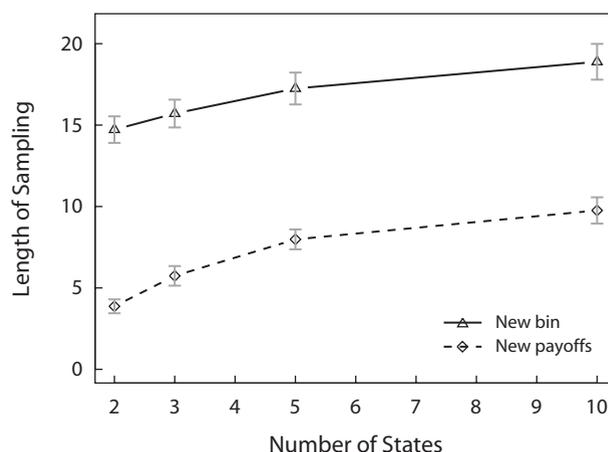


Figure 8. Mean length of sampling as a function of number of states and whether the bin composition changes (“New bin”) or only the payoffs in each state (“New payoffs”). Vertical bars indicate standard errors.

estimates of the variance depends). Often, there is a functional relationship between kurtosis and variance, such as with the Gaussian distribution; in those cases, effectively only variance is needed. These considerations lead to the following predictions:

- M1. Mean–variance optimizers resample even if the underlying state structure (number of colors and number of balls of each color) does not change.
- M2. Sampling length depends on moments of the payoff distribution, such as variance.

As was mentioned before, mean–variance optimizers ignore states and, as a result, may end up making statewise-dominated choices:

- M3. Mean–variance optimizers may purchase gambles at prices at least as high as the highest possible payoff or pass on gambles at a price at most equal to the lowest possible payoff.

Results. Twenty-seven subjects (10 women, 17 men) were recruited for the study. Twenty-four participants were students from the Swiss Federal Institute of Technology. The remaining 3 participants were scientific collaborators working at the Swiss Federal Institute of Technology or the University of Lausanne. The mean age was 24 years. The study was approved by the ethics committee of the Swiss Federal Institute of Technology. The decision-making task was programmed in E-Prime and run on location at the Laboratory for Decision Making Under Uncertainty. The mean pay-for-performance was 69 Frs, with a minimum of 47 Frs and a maximum of 91 Frs.

We observed very few cases in which the subjects chose statewise-dominated gambles (Hypothesis M3), suggesting that they paid attention to states and the payoffs associated with each state. On average, the subjects made 4.3% *buy* errors (buying a gamble that was dominated by its price) and 2.8% *pass* errors (passing on gambles when

the price equaled the lowest possible payoff). So, even if the subjects were mean–variance optimizers otherwise, they paid sufficient attention to states not to be frequently subject to the most damaging feature of this type of decision making, consistent with Hypothesis E3. Our findings call into question the economic relevance of statewise dominance mistakes reported in the literature (Mellers et al., 1995).

However, the probability for a subject to resample when we changed the labels (payoffs) of the balls in the bin, keeping the composition of the bin the same, was above 50% (dashed line in Figure 7). This indicates that the subjects partially relied on mean–variance optimization (Hypothesis M1). Yet the probability to resample also increased with the number of states. There is thus a surprising interaction with state complexity (number of states) that is not predicted by mean–variance optimization (M2) alone.

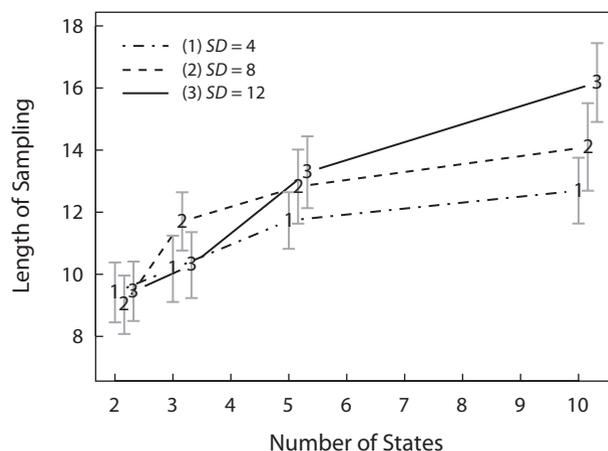


Figure 9. Mean length of sampling as a function of number of states, stratified by payoff standard deviation (SD). Vertical bars indicate standard errors.

In fact, when plotting the length of sampling against number of states, we observe an increase that suggests expected utility calculations (Hypothesis E2); see Figure 8 (length set to 0 if no sampling). As such, the data reject the clear-cut categorization into expected utility analysis or mean–variance analysis.

What is going on? It appears that the length of sampling also increases with the standard deviation of the payoff for 10 states and, to a lesser extent, for 5 states (see Figure 9). This increase is consistent with Hypothesis M2 but occurs only for a higher number of states.

Consequently, the overall picture that emerges is one in which subjects mix expected utility analysis (they increase sampling as state complexity increases; E2) and mean–variance analysis (they resample; M1), but once the number of states reaches a certain level, they tend to weigh mean–variance analysis more, increasing the length of resampling when payoff variance increases (M2).

Final Remarks

Mean–variance preference theory assumes that payoffs are evaluated in a linear way. That is, mean payoff is computed as the expectation of payoffs that are not transformed in a nonlinear way, as in standard expected utility theory. Likewise, payoff variance is just what it says: the mathematical variance of payoff, and not, for instance, the variance of the utility of payoff. Allais (1953) has suggested, however, that human choice under uncertainty can be represented better in terms of trade-offs of expected utility of reward and variance of reward utility. That is, expectations and variances are computed for nonlinear transformations of actual rewards and, therefore, on the basis of experienced utility. To date, it is not known whether encoding of expected rewards and reward variances in the human brain actually represent expected reward utility and reward utility variance. A new paradigm is needed to discriminate between standard mean–variance theory as a basis of choice and Allais's approach. Just like prospect theory preferences, Allais's preferences can accommodate most of the violations to the stringent rational choice axioms underlying standard expected utility theory that are commonly observed in human choice under uncertainty.

Here, we have focused on the mean–variance model of decision making. There are other models that separate mean and risk, whereby risk is measured in different ways. Most important, there is the mean-risk model with the coefficient of variation (the ratio of standard deviation over mean) as metric for risk. There exists evidence that the coefficient of variation is a better predictor of decision making under risk than is variance (or standard deviation) in both nonhuman primates (McCoy & Platt, 2005) and humans (Weber, Shafir, & Blais, 2004).

Our article has contrasted the mean–variance and expected utility approaches to modeling choice under uncertainty. Much of the discussion may seem to imply that we consider computation of values as an *either/or* situation: the brain encodes the features of gambles either in terms of probabilities and magnitudes (expected utility approach) or in terms of expected payoff and variance of payoff (mean–variance approach). One should keep an open mind, how-

ever, and envisage the possibility that the brain actually uses both approaches. As we discussed above, the mean–variance approach has the advantage of speed, certainly in a learning environment, but it is prone to mistakes. In contrast, learning in the expected utility approach is cumbersome but allows for powerful extrapolation when new gambles (defined on the same state space) are presented. It cannot be excluded that the human brain is capable of doing both, even simultaneously, emphasizing the output of one or the other, depending on the situation.

This suggestion amounts, of course, to an application of dual-system theory of cognition (Evans, 2003) to choice under uncertainty. It is analogous to other substantiations of dual-system theory, such as reflexive/reflective learning (Daw et al., 2005) or heuristic/logical problem solving (Kahneman & Frederick, 2002).

The experimental findings that we reported on here provided corroborating evidence for a dual system for decision making under uncertainty. Aspects of subjects' decisions supported expected utility, whereas other aspects indicated mean–variance analysis. We observed that when the number of states increased, subjects tended to rely more on mean–variance analysis. Because expected utility analysis is more demanding when the number of states is high, this change might be explained by the limited capacity of human information processing. In this context, it is interesting to note that we observed a significant effect of payoff standard deviation on sampling for 5 states; the effect became even more pronounced for 10 states. This may be explained by the fact that effective human working memory is limited to seven plus/minus two items (Miller, 1956). Further research is needed to provide explicit association of our result with working memory capacity.

AUTHOR NOTE

Correspondence concerning this article should be addressed to M. d'Acremont, Laboratory for Decision Making Under Uncertainty, Swiss Federal Institute of Technology, Odyssea, Station 5, CH-1015 Lausanne, Switzerland (e-mail: mathieu.dacremont@epfl.ch).

REFERENCES

- ALLAIS, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, **21**, 503-546.
- BELL, D. E. (1995). Risk, return, and utility. *Management Science*, **41**, 23-30.
- BIRNBAUM, M. H., & NAVARRETE, J. B. (1998). Testing rank- and sign-dependent utility theories: Violations of stochastic dominance and cumulative independence. *Journal of Risk & Uncertainty*, **17**, 49-78.
- BLACK, F., & SCHOLES, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, **81**, 637-654.
- BORCH, K. (1969). A note on uncertainty and indifference curves. *Review of Economic Studies*, **36**, 1-4.
- CHANDRASEKHAR, P. V. S., CAPRA, C. M., MOORE, S., NOUSSAIR, C., & BERNIS, G. S. (2008). Neurobiological regret and rejoice functions for aversive outcomes. *NeuroImage*, **39**, 1472-1484.
- D'ACREMONT, M., LU, Z.-L., LI, X., VAN DER LINDEN, M., & BECHARA, A. (2007, September). *Risk prediction in the human brain: A functional neuroimaging study*. Paper presented at the annual conference of the Society for Neuroeconomics, Boston.
- DAW, N. D., NIV, Y., & DAYAN, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, **8**, 1704-1711.

- DIACONIS, P., & FREEDMAN, D. (1986). On the consistency of Bayes estimates. *Annals of Statistics*, **14**, 1-26.
- DICKHAUT, J. (2008). *The emergence of economic institutions: A neuronal perspective* (Working paper). Minneapolis: University of Minnesota.
- EVANS, J. S. B. T. (2003). In two minds: Dual-process accounts of reasoning. *Trends in Cognitive Sciences*, **7**, 454-459.
- HAMPTON, A. N., BOSSAERTS, P., & O'DOHERTY, J. P. (2006). The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *Journal of Neuroscience*, **26**, 8360-8367.
- HARSANYI, J. C., & SELTEN, R. (1988). *A general theory of equilibrium selection in games*. Cambridge, MA: MIT Press.
- HUETTEL, S. A., SONG, A. W., & MCCARTHY, G. (2005). Decisions under uncertainty: Probabilistic context influences activation of prefrontal and parietal cortices. *Journal of Neuroscience*, **25**, 3304-3311.
- KAHNEMAN, D., & FREDERICK, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49-81). Cambridge: Cambridge University Press.
- KAHNEMAN, D., & TVERSKY, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47**, 263-291.
- KNUTSON, B., & BOSSAERTS, P. (2007). Neural antecedents of financial decisions. *Journal of Neuroscience*, **27**, 8174-8177.
- KNUTSON, B., TAYLOR, J., KAUFMAN, M., PETERSON, R., & GLOVER, G. (2005). Distributed neural representation of expected value. *Journal of Neuroscience*, **25**, 4806-4812.
- MARKOWITZ, H. (1952). Portfolio selection. *Journal of Finance*, **7**, 77-91.
- MCCLURE, S. M., BERNS, G. S., & MONTAGUE, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, **38**, 339-346.
- MCCOY, A. N., & PLATT, M. L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nature Neuroscience*, **8**, 1220-1227.
- MELLERS, B. A., BERRETTY, P. M., & BIRNBAUM, M. H. (1995). Dominance violations in judged prices of two- and three-outcome gambles. *Journal of Behavioral Decision Making*, **8**, 201-216.
- MILLER, G. A. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, **63**, 81-97.
- O'DOHERTY, J., DAYAN, P., SCHULTZ, J., DEICHMANN, R., FRISTON, K., & DOLAN, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*, **304**, 452-454.
- OGRYCZAK, W., & RUSZCZYŃSKI, A. (1999). From stochastic dominance to mean-risk models: Semideviations as risk measures. *European Journal of Operational Research*, **116**, 33-50.
- PAULUS, M. P., ROGALSKY, C., SIMMONS, A., FEINSTEIN, J. S., & STEIN, M. B. (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., BOSSAERTS, P., & QUARTZ, S. R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron*, **51**, 381-390.
- PREUSCHOFF, K., QUARTZ, S. R., & BOSSAERTS, P. (2008). Human insula activation reflects risk prediction errors as well as risk. *Journal of Neuroscience*, **28**, 2745-2752.
- RESCORLA, R., & WAGNER, A. (1972). A theory of classical conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. Black & W. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- ROLLS, E. T., MCCABE, C., & REDOUTE, J. (2008). Expected value, reward outcome, and temporal difference error representations in a probabilistic decision task. *Cerebral Cortex*, **18**, 652-663.
- SAVAGE, L. J. (1972). *The foundations of statistics*. New York: Courier Dover.
- TOBLER, P. N., FIORILLO, C. D., & SCHULTZ, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science*, **307**, 1642-1645.
- TOBLER, P. N., O'DOHERTY, J. P., DOLAN, R. J., & SCHULTZ, W. (2007). Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *Journal of Neurophysiology*, **97**, 1621-1632.
- TOM, S. M., FOX, C. R., TREPEL, C., & POLDRACK, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, **315**, 515-518.
- TVERSKY, A., & KAHNEMAN, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk & Uncertainty*, **5**, 297-323.
- VON NEUMANN, J., & MORGENSTERN, O. (1947). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- WEBER, E. U., SHAFIR, S., & BLAIS, A. R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, **111**, 430-445.
- YACUBIAN, J., GLÄSCHER, J., SCHROEDER, K., SOMMER, T., BRAUS, D. F., & BÜCHEL, C. (2006). Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *Journal of Neuroscience*, **26**, 9530-9537.
- YACUBIAN, J., SOMMER, T., SCHROEDER, K., GLÄSCHER, J., BRAUS, D. F., & BÜCHEL, C. (2007). Subregions of the ventral striatum show preferential coding of reward magnitude and probability. *NeuroImage*, **38**, 557-563.

NOTE

1. Only in prospect theory do belief biases contribute to risk attitudes.

APPENDIX

Formal Definitions of Expected Utility and Mean-Risk Models

Here, we will provide formal definitions of expected utility and mean-risk modeling. We will also show how the two models are related, by means of Taylor series expansions. In expected utility (EU) theory, the value V of an option is defined as the expectation of the utility of the payoff x across states:

$$V_{\text{EU}} = E[u(x)] = \sum_{s=1}^n p_s u(x_s), \quad (1)$$

where n is the number of states, p_s the probability of state s , and x_s the return associated to state s .

Prospect theory (PT) is an amendment to this basic approach, whereby probabilities are also transformed:

$$V_{\text{PT}} = \sum_{s=1}^n w(p_s) u(x_s), \quad (2)$$

where w is the probability-weighting function. (Unlike in expected utility theory, u is usually called the *value* function in prospect theory.) In standard expected utility theory, risk attitudes are gener-

APPENDIX (Continued)

ated through nonlinearities in u . Specifically, if u is strictly concave (meaning that it increases at a decreasing rate), the resulting choices under uncertainty will reflect risk aversion. In prospect theory, risk sensitivity results from a combination of nonlinearities in the value function and the presence of a probability-weighting function.

In mean-risk (MR) models, the value of a risky option is expressed as a trade-off between a return and risk measure (Bell, 1995):

$$V_{MR} = r(x) - R(x), \quad (3)$$

where r denotes the return function and R the risk function.

In certain conditions, the two models become equivalent. For instance, if the utility function is quadratic,

$$u(x) = ax^2 + bx + c, \quad (4)$$

the expected utility can be written as a mean-risk model:

$$V_{EU} = E[u(x)] = u(\mu) + a\sigma^2, \quad (5)$$

where μ and σ^2 are the expected value and variance of the payoff x .

In general, utility functions can always be approximated by means of a Taylor series expansion:

$$u(x) \cong u(\mu) + u'(\mu)(x - \mu) + \frac{1}{2}u''(\mu)(x - \mu)^2, \quad (6)$$

where u' and u'' denote the first and second derivatives of u . Taking the expected value of the utility, we obtain

$$V_{EU} \cong u(\mu) + \frac{1}{2}u''(\mu)\sigma^2, \quad (7)$$

which has the form of a mean-risk model. As such, valuations from expected utility theory can always be approximated with valuations from the mean-risk approach.

The mean–variance model is a special case of the mean-risk model. In it, values are computed as follows:

$$V_{MR} = \mu - b\sigma^2, \quad (8)$$

where b reflects a penalty for risk, and hence, if $b > 0$, values will generate risk avoidance. This mean–variance model is the valuation approach of choice in the area of finance.