Markowitz in the brain ?

par Kerstin PREUSCHOFF, Steven QUARTZ et Peter BOSSAERTS

| Dalloz | Revue d'économie politique

2008/1 - Volume 118
ISSN 0373-2630 | pages 75 à 95

Pour citer cet article :
— Preuschoff K., Quartz S. et Bossaerts P., Markowitz in the brain ?, Revue d'économie politique 2008/1, Volume 118, p. 75-95.

Distribution électronique Cairn pour Dalloz.

Markowitz in the brain ?

Kerstin Preuschoff*
Steven Quartz**
Peter Bossaerts***

We review recent brain-scanning (fMRI) evidence that activity in certain sub-cortical structures of the human brain correlate with changes in expected reward, as well as with risk. Risk is measured by variance of payoff, as in Markowitz' theory. The brain structures form part of the dopamine system. This system had been known to regulate learning of expected rewards. New data show that it is also involved in perception, of expected reward, and of risk. The findings suggest that the brain may perform a higher-dimensional analysis of risky gambles, as in standard portfolio theory, whereby risk and expected reward are considered separately. That is, the human brain appears to literally record the very inputs that have become a defining part of modern finance theory.

neurofinance - neuroeconomics - decision making under uncertainty - Markowitz - portfolio theory - dopaminergic - system

Markowitz dans les cerveaux ?

Nous décrivons les résultats récemment obtenus par moyen de fMRI sur la corrélation d'activations de certaines structures sub-corticales avec les changements d'espérance et de risque de revenus monétaires. Nous utilisons l'écart-type pour mesurer le risque. L'ecart-type constitue une mesure de risque qui correspond à celle utilisée dans la théorie de Markowitz. Les structures cérébrales que nous étudions font partie du système dopaminergique. Ce système était déjà impliqué dans l'apprentissage d'espérances de gains. Nos données indiquent que ce système est également impliqué dans la perception, aussi bien d'espérance que de risque. Nos résultats suggèrent que les cerveaux de l'homme appliquent une analyse multi-dimensionelle de jeux aléatoires, comme dans la théorie traditionnelle de gestion de portefeuilles. Là aussi, on considère séparément l'espérance et le risque. Ainsi, le cerveau humain semble littéralement encoder des mesures qui ont partie intégrante de la théorie des finances modernes.

neurofinance - neuroéconomie - prise decision dans l'incertitude - Markowitz – théorie de gestion de portefeuilles - système dopaminergique

Classification JEL : D87, D81, G11, C91

* University of Zurich
** Caltech
*** EPFL and Caltech (corresponding author)

Neoclassical Finance treats the human agent as a black box. It assumes that the agent takes in information and converts it into actions "as if" she were a computer that blindly applies certain rules. The rules are those of rational decision making. Neoclassical Finance is silent about how exactly data in the environment are perceived and how perception is translated into choice. As a positive theory of human decision making, this approach has really never worked. Starting with Allais [1953], numerous instances have been cited where most people make choices that are at odds with the theory.

In contrast to Neoclassical Finance, *Behavioral Finance* starts from observation of how investors actually choose (or from surveys about how they would choose). Like Neoclassical Finance, however, Behavioral Finance still treats the human agent as a black box. Choices are made "as if" the agent were a computer that invokes certain rules. In the case of Behavioral Finance, these rules do not come from any normative theory of decision making. Instead, the rules are inferred from extensive observation of actual choices in various contexts, or through many questionnaires that probe agents' likely choices in numerous hypothetical situations.

Human behavior is, however, bewilderingly complex and heterogeneous. The list of cognitive biases (deviations of actual behavior from rational choice) is long (see, e.g., Pohl [2005]) and, given the context-dependency ("framing") of many biases, likely to lengthen substantially. Moreover, it seems rather superficial to determine a categorization of human choice purely on readily observable or recordable features (behavior or surveys). Altogether, the endeavor risks to become as painful and inaccurate as categorization of plant and animal life on the basis of external features was before biologists discovered DNA. For that matter, the discovery of DNA was not an outgrowth of biologists' efforts to categorize plant and animal life, but came about only because biologists started looking beyond the external features, and into the inner workings of plants and animals. In our view, only an analogous change in methodology may lead to the revolution of financial theory that Behavioral Finance is looking for.

This paper discusses results from an alternative approach to studying human decision making under uncertainty. Our approach focuses directly on the box where perception and choice are "done:" the brain. The purpose is to explore the inner workings of the brain, with the hope of eventually coming to a better understanding of externally observed cognition.

The particular question that concerns us here is whether risks and rewards are encoded in the brain in terms of a single-dimensional index, or whether they are somehow decomposed. With few exceptions, ever since Bernoulli proposed logarithmic expected utility, decision theorists and economists have represented the features of risky gambles in terms of a one-dimensional index. This includes Prospect Theory, which is basically a representation of choice (hypothetical or actual) in terms of a subjective expected utility metric.

Finance, however, has always found it useful to decompose this index using a Taylor series expansion argument, so that the different facets of risk (expectation, variance, skewness, kurtosis, etc.) become transparent. An agent with logarithmic expected utility, for instance, prefers higher expected

return, is averse to variance, prefers positive skewness (while averse to negative skewness), and is averse to kurtosis.

In fact, finance has discarded the idea of a single-dimensional utility index altogether, since risky gambles became evaluated by contrasting the scores they get along the various statistical dimensions. In the earliest versions of this approach, only the first two moments (expectation and variance) were traded off against each other (Markowitz [1952]). The approach is extremely useful. Among other things, analysis of several simultaneous gambles (portfolio analysis) is vastly simplified in this "Markowitz" world. Moreover, the approach facilitates learning. As evidence accumulates, gambles can be re-evaluated by updating the various moments, a rather straightforward computational exercise.

Because of these advantages, we conjectured that the human brain actually followed the finance approach, encoding the various statistical inputs needed for effective evaluation of the desirability of risky gambles. As we shall discuss later, close inspection of recent evidence from the non-human primate brain (Fiorillo, Tobler and Schultz [2003]) as well as from the human brain (Knutson, e.a. [2003]) strengthened our belief that the conjecture was right.

As we shall discuss, the human brain appears to literally record the very inputs that have become a defining part of modern finance theory. In particular, neurons in parts of the brain respond immediately (with minimal delay) to changes in expected reward and with a short delay (about 1 to 2 seconds), to risk, as measured by payoff variance.

Admittedly, variance is only a partial measure of risk. At the level of risk that is typical in laboratory experiments, however, it has so far been hard to find evidence of higher-order risk (skewness aversion, for instance). Prices in large-scale laboratory financial markets, for instance, conform with the Capital Asset Pricing Model (CAPM), as if only variance matters (Bossaerts and Plott [2004]). Although some may be surprised to find evidence of risk aversion at all, the phenomenon is a widely documented and easily replicable one. See Holt and Laury [2002] for a detailed analysis.

The findings discussed here were first reported in the neuroscience literature (Preuschoff, Bossaerts and Quartz [2006]). The goal of this paper is to explain the findings to a finance audience. To do this effectively, the paper includes a brief overview of relevant facts about the human and non-human primate brain. The remainder of the paper is organized as follows. Section I introduces neuroscience terminology and explains how brain activity can be measured. Section II surveys earlier results about reward expectation as recorded in primate and human brains. Our own results are presented in Section III. Section IV provides further perspective.

I. The Brain And Measuring Brain Activity

This section introduces basic terminology about brain activity. It also introduces some brain regions of interest for decision making under uncertainty, as well as two methods that neuroscientists use to better understand processes in the brain.

Modern neuroscience is built on the foundations the Spanish anatomist Santiago Ramón y Cajal articulated in the late nineteenth century. Known as the Neuron Doctrine (see, e.g., Bullock, e.a. [2005]), Cajal brilliantly hypothesized on the basis of his anatomical observations that brain cells, neurons (which number on the order of 10 11 in the human brain) are polarized so that signals flow unidirectionally, as schematized in Figure 1, from receptive structures, dendrites, to the cell body (soma), where they are integrated and then outputted along axons to other neurons. A defining tenet of the Neuron Doctrine, for which Cajal shared the Nobel Prize in 1906, was that neurons are individual functional units that communicate with one another across synapses (Fig 1 inset).

Modern neuroscience is also built on the principle that the signal processing power of neurons rests on how incoming signals are nonlinearly transformed into outputs. The key notion is that incoming signals must reach a threshold in order for the neuron to transmit a spike of electrical activity, or action potential. Early on, it was thought that action potentials amounted to a binary code of all-or-nothing spikes. Indeed, in 1945 John von Neumann pointed to this feature of neurons as justification for the digital design of modern computers. Much of neuroscience research focuses on measuring and characterizing these spikes, which turn out to be far more nuanced than originally proposed and constitute the neural code that underlies our sensations and actions.

Unlike the electrical relays of modern computer circuits, the electrical impulses of neuronal spikes are converted into chemical signals when neurons communicate across synapses. When a spike reaches an axon terminal, a small gap, the synaptic cleft, is bridged by means of chemicals called *neurotransmitters,* whose effect is determined by specialized receptors on the postsynaptic neuron. Excitatory neurotransmitters act like currency in an economy, because they facilitate transactions by making action potentials more likely, while inhibitory neurotransmitters are more like debt collectors, because they make action potentials less likely In recent years, neuroscientists have increasingly appreciated the bewildering complexity of neuronal signal processing: there are hundreds of different neurotransmitters and receptors, acting at many different timescales, and neurons typically receive input from thousands of other neurons.

One neurotransmitter we will discuss at length is *dopamine,* which is released by *dopaminergic neurons* of the midbrain, in the *ventral tegmental area* and the *substantia nigra* (Fig 2). Although these cells constitute less than 1% of the neurons in the brain, they have a profound effect on brain

function and are involved in a wide array of functions, including voluntary movement, mood and arousal, and reward, and they have been implicated in Parkinson's disease, attention deficit hyperactivity disorder, Tourette's syndrome, schizophrenia, bipolar disorder, and most forms of drug abuse.

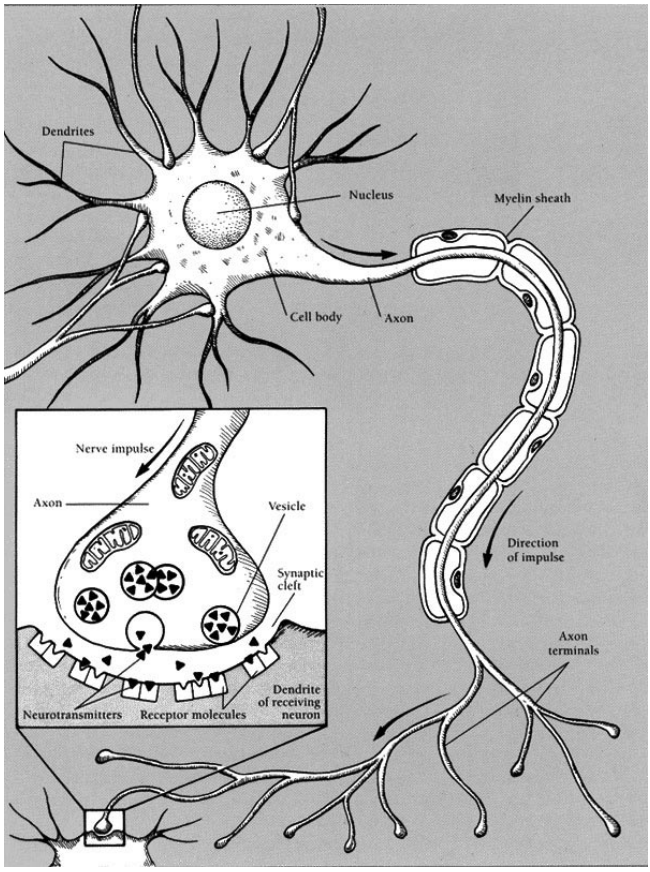

Figure 1. Sketch of a neuron (From: *Brain Facts,* Society for Neuroscience [2004])

Such a relatively small number of dopaminergic neurons can exert profound effects on brain function because of their anatomical organization. Dopaminergic neurons reside in subcortical structures buried deep inside the brain and project up into other subcortical structures and to the cortex, a thin, folded six-layered sheet of neurons, whose folds give the brain its characteristic appearance. As Figure 2 displays, dopaminergic neurons in the Ventral Tegmental Area (VTA) and in Substantia Nigra (SN) have extremely widespread projection areas, or targets. These projection regions

are collectively known as *dopaminoceptive regions.* In the sequel, we will focus on one projection region, namely, the *nucleus accumbens (ventral striatum).* The messages that dopaminergic neurons send to this region have been shown to reflect, among other things, prediction errors. As such, dopamine plays a crucial role in *learning.*

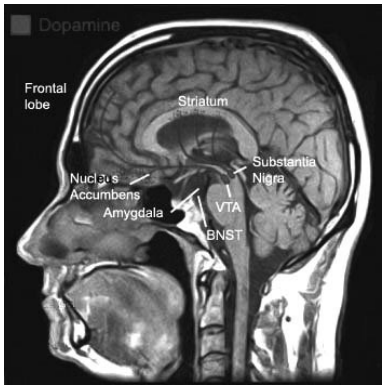

Figure 2. Sketch of part of the dopamine system in the brain.

Neuroscientists study brain activity in many ways. Two procedures are of particular interest to us. Single-unit electrophysiological recordings directly monitor activation of single neurons. The other procedure, functional magnetic resonance imaging (fMRI), indirectly monitors the activity of groups of neurons, through their effect on oxygen levels in the surrounding bloodstream.

In single-unit electrophysiology recordings, microelectrodes are placed in a (mostly non-human) brain to reach neurons in a region of interest and record spikes from single neurons. Changes in the number of spikes over time reflect changes in neuronal activity. For instance, a dopaminergic neuron will increase the number of spikes over time in response to an unexpected reward.

While electrophysiology has both high temporal and spatial resolution, its invasive nature makes it generally prohibitive in humans, except for limited cases of neurosurgery patients. Until relatively recently, this meant that neuroscience relied mainly on animal models of human brain function. With the advent of functional magnetic resonance imaging (fMRI), neuroscience has made significant progress in investigating human brain function directly. fMRI detects changes in blood oxygen levels in the brain using a strong magnetic field and is based on the idea that any neural activity requires energy. Since the brain is almost entirely dependent on aerobic forms of energy production, brain activity is tightly linked to oxygen delivery through blood flow. The more active a region, the more oxygen it requires. Thus the blood oxygen level dependent, short BOLD, changes with activity. fMRI is an indirect measure of brain activity because it does not measure spikes but the increased blood flow that is the consequence of an increase in spiking.

For example, dopaminergic neurons spike more frequently in order to convey larger reward anticipation. The increase in spiking induces activation of neurons in nucleus accumbens. The activation is only possible, however, with more energy, and hence, increased oxygen-rich blood flow. As a result, positive correlation will emerge between a quantity of interest to finance, namely, reward anticipation, and brain activation, measured in terms of BOLD signals. The correlation allows us to ask additional questions, such as whether BOLD changes reflect changes in reward anticipation when the latter is measured as in finance, namely, as the (mathematical) expectation of the payoff.

II. Expected Reward And Risk Perception

For a while now, it has been known that the activity of dopaminergic neurons correlates nicely with changes in reward anticipation. More recently, a link between risk and spiking of these neurons has been discovered. We first describe the evidence of the former, in non-human primates as well as in humans. Subsequently, we discuss evidence for correlation with risk in the non-human primate brain.

• Reward Expectation

In a typical experiment, an animal is trained to associate a neutral stimulus (e.g., a picture), the *conditioning stimulus,* with a potential reward (e.g., juice) occurring a fixed amount of time after the stimulus. This is known as *Pavlovian conditioning.* Once the association has been learned the formerly neutral stimulus is still not rewarding in itself but predicts a reward. This reward may be stochastic both in its size and occurrence, in which case one could describe it as a gamble. Multiple associations can be learned simultaneously by using different stimuli (i.e. different pictures) for different reward magnitudes and probabilities.

Neuroscientists are interested in brain activation when the conditioning stimulus (picture) is presented to a fully-conditioned animal, and how brain activation changes as a function of the stimulus, and hence, as a function of the nature of the gamble.

In particular, neuroscientists discovered that the intensity of the spike trains of some dopaminergic neurons correlated with *reward expectation:* when the conditioning stimulus arrived, the spiking of certain neurons increased with the probability of reward delivery. In addition, at the time of reward, spiking of these neurons correlated with the prediction error – the difference between the actual reward (juice squirt or not) and the conditional expectation of the reward.

This is depicted in Figure 3, reproduced from Fiorillo, Tobler and Schultz [2003]. Figure 3 shows single-neuron spike trains for several trials of each of five trial types. Each trial type corresponds to a different gamble. Across the gambles, the probability of a (fixed-size) juice squirt differed, from zero to one. The actual visual stimulus that cued the reward probability p in each trial type is displayed below the numerical value of the probability. Figure 3 only shows the trials for which the reward did occur. CS indicates the time of the conditioning stimulus; *Reward* indicates the time of the reward delivery. Each sequence of vertical dashes under the horizontal axis corresponds to the spike train of one trial. The spikes are aggregated into small bins over time and then averaged across trials. This average is displayed above the horizontal line.

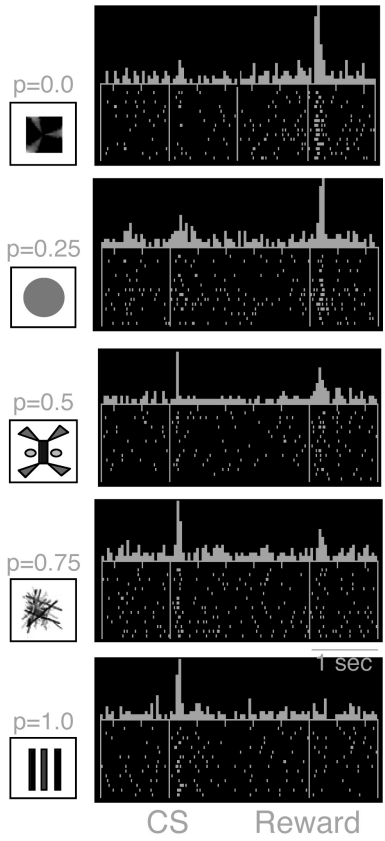

Figure 3. Spike trains for a single neuron in the *ventral tegmental area,* for different gambles distinguished by probability of reward ; only rewarded trials are shown.
Source : Fiorillo, Tobler and Schultz [2003].

Figure 3 demonstrates that neuronal activity is correlated with the reward prediction error. For $p = 0.25$, the probability of spiking at the time of reward and conditional on reward delivery is about 3 times as large as that at the time of the stimulus. Conversely, for $p = 0.75$, the probability of spiking at the time of reward and conditional on reward delivery is only one third as large as at the time of the stimulus. At $p = 0.5$, the relative magnitudes are not exactly equal as expected, but this is an effect of the (arbitrary) bin size. When $p = 0$, reward delivery is theoretically impossible. Figure 3 displays neuronal activity in the few trials where, contrary the monkey's expectation, reward was delivered. The neuronal response is commensurate.

An analogous phenomenon has been found in human studies. For example, Knutson, e.a. [2003] used fMRI to study activation in, among others, some projection areas of the dopaminergic neurons, such as the *nucleus accumbens* in the *ventral striatum.* Stimuli (cues) for different gambles were shown. Each gamble led to a different reward. Reward delivery depended on whether the subject could push a button in a timely fashion. Otherwise the subject gained nothing. The button press was made difficult enough so that subjects succeeded only 66% of the time. That is, the gambles paid with a chance equal to 0.66. A few seconds later, the outcome was displayed. Rewards accumulated across trials.

Figure 4 shows time courses of the average fMRI signal in the nucleus accumbens ("NAcc" in the Figure) and 95% confidence intervals (vertical line segments) for each gamble. Gambles are distinguished by reward (upon successful button press). The white regions indicate (i) the time period during which the subject learns about the nature of the gamble ("ant"), (ii) the time period when the subject is informed about the outcome ("out").[1] Average time courses are shown only for trials that were rewarded. Notice how the average fMRI signal increases significantly in the expected reward upon presentation of the stimulus. This is consistent with the idea that dopaminergic neurons fire more when reward expectation increases. Likewise, activation increases significantly in the reward level at the time the reward is announced.

Unfortunately, since the subjects themselves generated the uncertainty, it is not clear whether the observed activation is purely perceptual instead of motivational. When the stimulus indicated a higher expected reward, a subject may have been more *motivated* to push the button than when the stimulus indicated a lower expected reward, and the observed activation could merely reflect this motivation. We will confirm, however, that brain activation emerges without motivation. That is, the activation is perceptual.

The time resolution of the fMRI signal (displayed in Figure 4) is not as good as that of electrophysiological measurement of activity of single neurons (compare to Figure 3). As a result, it may seem that there is already a response *before* the stimulus is presented, as if the subject anticipates the type of gamble (s)he is about to play. We know from single-neuron mea-

1. The stimulus as well as the outcome were shown approximately 4 seconds earlier. The 4 second gap captures typical delay for the signal to reach its peak and hence to be captured by fMRI.

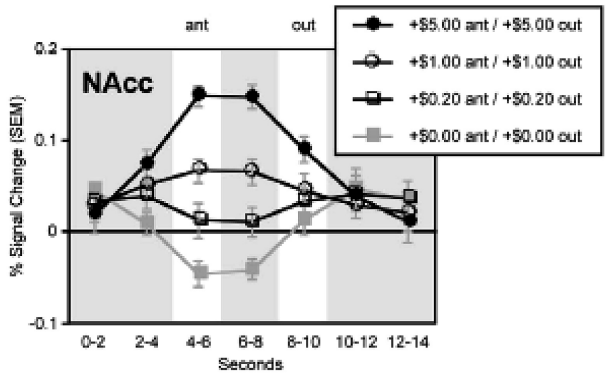

Figure 4. Average percentage change in fMRI signal over time, as a function of expected reward – rewarded trials only. Vertical line segments indicated 95 % confidence intervals. « ant » is period during which stimulus is expected to have a measurable fMRI response (4" after actual presentation of stimulus) ; « out » is period during which outcome revelation is expected to have a measurable hemodynamic response. Source : Knutson e.a. [2003].

surements (see Figure 3) that this is not the case. The anomaly is just a consequence of lower time resolution.

In addition, there seems to be a sustained response to expected reward *after* the stimulus is recorded (in the grey area between "ant" and "out" in Figure 4). Again, this would be inconsistent with the evidence from single-neuron recordings (see Figure 3). While one may be inclined to again attribute the anomaly to lower time resolution of the fMRI signal, there is in fact a more plausible explanation: risk perception.

As we discuss next, some dopaminergic neurons do not respond to expected reward, but respond to risk as measured by variance, albeit with a delay. The firing of some dopaminergic neurons in response to risk may translate into a (delayed) fMRI signal in the projection areas, such as, in this case, the nucleus accumbens. In the Knutson e.a. [2003] setup, however, one cannot distinguish between this hypothesis and an explanation based on low time resolution of the fMRI signal, because risk increases with expected reward. Instead, we need to vary risk independently of expected reward and check whether the sustained response to expected reward disappears, yet increases with risk. This was the goal of our own experiment, to be discussed in Section IV.

Risk Perception

Let us momentarily turn back to the Fiorillo, Tobler and Schultz [2003] study of neuronal spiking in the *ventral tegmental area* (VTA) of the monkey

brain in reaction to conditioning stimuli that signal different gambles. Remember that gambles are differentiated by the probability p that a (fixed) juice squirt is delivered. In their setting, p was any of the following: 0, 0.25, 0.5, 0.75, 1. As p changes, not only does expected reward change, but also *risk* as measured by variance. Risk is maximal when $p = 0.5$; it is minimal when either $p = 0$ or 1. Fiorillo and his collaborators noticed that certain neurons in the VTA respond in a delayed, but sustained fashion to risk. These neurons turn out to largely a different population from those that react immediately to expected reward. The sustained response takes on an intriguing form: when averaged across neurons, the response increases with time.

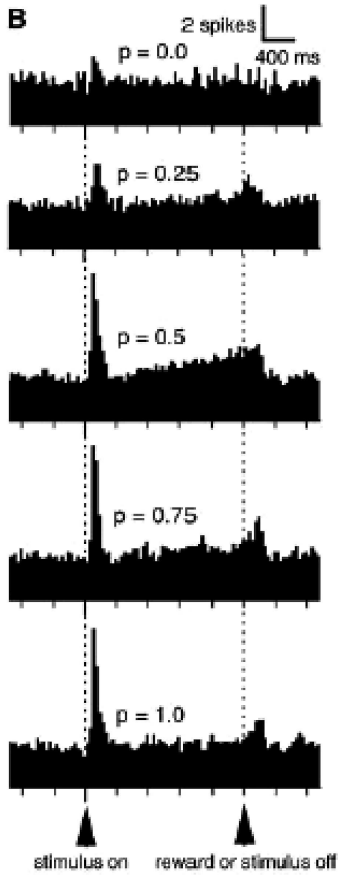

Figure 5. Average spike trains across 35 to 44 neurons in the ventral tegmental area, for different gambles distinguished by probability of reward ; all trials are included (both rewarded and unrewarded). Source : Fiorillo, Tobler and Schultz [2003].

Figure 5 provides the evidence. In it, the firing of all neurons under investigation is averaged, across both rewarded and unrewarded trials. Notice that the delayed response appears to be increasing in risk; it is maximal when risk is maximal (*i.e.*, when $p = 0.5$).

The immediate response of neurons that react to expected reward at the moment the conditioning stimulus presentation is also evident in Figure 5. This response increases with probability, and hence, with expected reward.

It may be surprising that we also observe a short-term response at the moment risk is realized, i.e., when the stimulus is switched off and the outcome (reward/no reward) is revealed. One would not have expected this because both rewarded and unrewarded trials are averaged over. In other words, the average prediction error should be zero. Yet VTA neurons react positively on average to the realization of the risk. It could be assumed that this is because of a fundamental asymmetry: neurons can reduce their firing rate (in response to absence of reward) only to zero; they can increase their firing rate (in response to reward delivery) to a much higher extent – at least in principle. It appears, however, that the recorded average is just a continuation of the delayed activation in response to risk (see Fiorillo, Tobler and Schultz [2005]).

In short, dopaminergic neurons appear to react to expected reward without delay, whereas they respond to risk with some delay, but in a sustained way. Let us now show that this phenomenon can be detected in fMRI analysis of the human brain as well.

III. Risk Perception And Sustained Neuronal Activity In The Ventral Striatum

• The Experiment

We scanned subjects while they repeatedly played the following simple card game. Ten cards numbered 1 to 10 where randomly shuffled. Two cards are to be drawn consecutively (and without replacement) from this deck. The subject is asked to guess whether the first or the second card is going to be the highest. Subjects place their bet. About 3 seconds later, the first card is displayed. About 7 seconds later, the second card is displayed, from which the subject can immediately infer whether (s)he has won. The subject is then asked to confirm whether (s)he won or lost. Twenty-five seconds after the beginning of this trial, a new trial starts. The setup is represented graphically in Figure 6.

Subjects played 3 sessions with 30 trials per session. Subjects start out each session with $25. They earn $1 if they guessed the right card. They lost

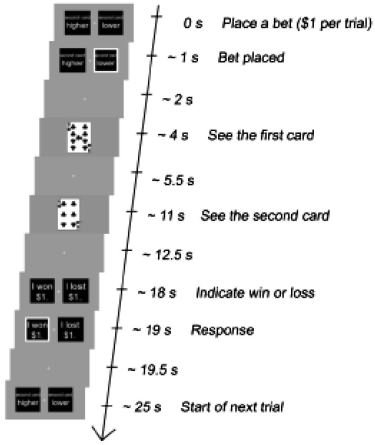

Figure 6. Timeline for the card experiment. To the left, screenshots are displayed that the subject saw at each stage of a trial. Source : Preuschoff, Bossaerts and Quartz [2006].

$1 if they were wrong. If no bet was placed, they lost automatically. They also lost $0.25 if they incorrectly indicated whether they had won/lost or if they did not respond. At the end of the experiment, the subject selected one of the three sessions at random which determined their final payoff.

When subjects place their bet, they have a 50% chance of winning. Since they are given $25 to play 30 gambles, they make approximately $0.83 per gamble plus the earnings from the gamble. At a 50% chance of winning, the expected payoff per gamble is $(0.83+0)=$0.83 (we assume that they always correctly indicate whether they win/lose). The risk (standard deviation) per gamble is $1.00. When the first card is shown, subjects received a *cue* that both the expected reward and the risk have changed. If the subject bets that the second card is going to be higher, and the first card is 8, then the expected reward is lowered to $(0.83-0.56) = $0.27, but the risk is also reduced, to $0.83. If the subject bets that the first card is going to be the higher of the two, then the expected reward is increased to $(0.83+0.56)=$1.39; the risk is again reduced to $0.83. The revised (conditional) expected rewards are linear in the revised (conditional) probabilities of winning; when measured as variance (instead of standard deviation), the revised risk is quadratic in the revised probabilities – with a maximum at a 50% chance of winning.

• The Results, Part 1 : Expected Reward

To determine whether there was any relationship between neuronal activity in the *nucleus accumbens* and expected reward, we project the fMRI signal in a cluster of recordings (voxels) onto dummy variables, one for each

level of probability of winning conditional on the cue (first card). This projection is done per subject. The results are corrected for serial correlation in the errors. Mathematically, let $y(i,t)$denote the (de-trended, normalized) fMRI signal in a region of interest (cluster of voxels) for subject i at time t. When the first card is shown, there are 10 levels of the (conditional) probability of winning: $p(j), j = 1,...10$. Let $X(j,t)$ denote a dummy variable that equals 1 for one second after the first card is presented and if the (conditional) probability of winning is $p(j)$. Let $Z(j,t)$ denote a dummy variable that equals 1 from one second after the first card is presented until the second card is shown and if the (conditional) probability of winning is $p(j)$. Let $W(t)$ denote a vector of (de-meaned) regressors that control for other events during a trial (visual cues and motor activity). We run the following regression:

$$y = (i,t) = \Sigma_j \beta(i,j) X(j,t) + \Sigma_j \beta(i,j+10) Z(j,t) + \pi \cdot W(i,t) + e(i,t).$$

We refer to the slope coefficients of the dummy variables as *betas.*

We obtain one vector of betas for each subject. Figure 7 displays the average (across subjects) $\beta(i,j)$ for $j = 1,...,10$, *i.e.*, the average estimated betas for the $X(j,t)$ dummy variables. Also shown are 95% confidence intervals based on the standard error of the estimated betas (averaged across subjects). Figure 7 summarizes the results for the left and right nucleus accumbens. It is clear that beta increases in the (conditional) probability of winning. That is, the fMRI signal increases in the (conditional) expected reward. Further tests (not reported) suggest that one cannot reject the null that the relationship is linear.

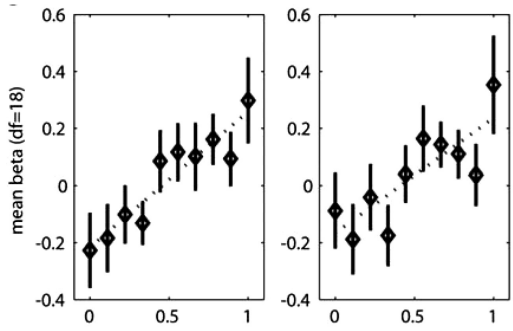

Figure 7. Relationship between mean betas and probability of winning in the left nucleus accumbens (left panel) and the right nucleus accumbens (right panel). Betas are slope coefficients for dummy variables that are set equal to one for a one-second period after cue presentation. Dummy variables vary with level of probability of winning. Vertical bars indicate 95% confidence intervals (df=degrees of freedom). Source : Preuschoff, Bossaerts and Quartz [2006].

Figure 8 provides a cross-section of the brain through the *ventral striatum.* Depicted are the regions that generate a significantly higher fMRI signal for cues that signal a high (conditional) expected reward as compared to cues that signal a low expected reward. The right and left nucleus accumbens are circled. The picture is obtained with the following projection. (This projection was also used to determine the size of the cluster in the projections described earlier.) Let $p(i,t)$ denote the probability of reward in the trial that subject i was given at time t. Let $E[x]$ denote the mean of the random variable x across trials throughout the entire session. Let $P(i,t)$ be a dummy variable which equals 1 for one second after display of the first card. Let $S(i,t)$ be a dummy variable which equals 1 from one second after display of the first card until presentation of the second card. We perform the following regression:

$$y(i,t) = \mu_1 [p(i,t) - 0.5]P(i,t) + \mu_2 \{ [p(i,t) - Ep(i)]^2 - E[p(i,t) - Ep(i)]^2 \} P(i,t) \quad [*]$$
$$+ \mu_3 [p(i,t) - 0.5]S(i,t) + \mu_4 \{ [p(i,t) - Ep(i)]^2 - E[p(i,t) - Ep(i)]^2 \} S(i,t)$$
$$+ \pi \cdot W(i,t) + e(i,t).$$

Figure 8 is based on the F-test for the coefficient μ_1. Therefore, the test checks whether there is a *linear* relationship in the *phasic* response of the fMRI signal to the conditional probability of reward.

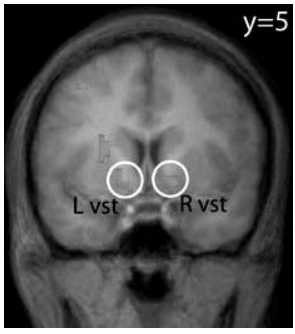

Figure 8. Coronal cross-section of the brain through the nucleus accumbens (ventral striatum ; vst) with regions of significantly higher activation for high vs. low conditional probability of winning upon display of the first card. L=left ; R=right. Source : Preuschoff, Bossaerts and Quartz [2006].

• The Results, Part 2: Risk

We next look at betas for dummy variables that are set equal to one for an interval starting one second after cue arrival and until card 2 is revealed

($Z(j,t)$ in the above equation). These dummies should capture the delayed neuronal response and its relationship with probability levels. Figure 9 shows the results, again for the *nucleus accumbens.*

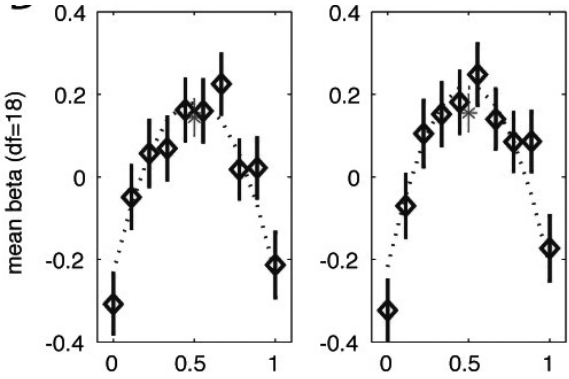

Figure 9. Relationship between mean betas and probability of winning in the left nucleus accumbens (left panel) and the right nucleus accumbens (right panel). Betas are slope coefficients for dummy variables that are set equal to one for period starting one seconds after presentation of the first card ("cue") and running till presentation of the second card. Dummy variables vary with level of probability of winning. Star indicates activation level before card 1 is shown, when the probability of winning equals 0.5. Vertical bars indicate 95% confidence intervals (df = degrees of freedom). Source : Preuschoff, Bossaerts and Quartz [2006].

The relationship between betas and probability levels is quadratic. Betas are maximal for intermediate probability levels and lowest for high or low levels. In other words, betas are approximately linearly related to risk (as measured by variance, which is quadratic in probability).

We implemented a simple consistency check, as follows. Before display of card 1, the probability of winning is 0.5. If the activation documented in Figure 9 indeed reflects risk, the level of activation before display of card 1 should correspond to the level recorded in Figure 9 for $p = 0.5$. The star indicates the average activation level in the same areas of the brain before display of card 1. The average is insignificantly different from the one for activation after display of card 1 when the number on card 1 implies that p remained at 0.5.

The evidence in Figure 9 fits with the finding in Fiorillo, Tobler and Schultz [2003] hat dopaminergic neurons in the ventral tegmental area react in a delayed fashion to risk. These neurons provide input to (project onto) neurons in, among others, the nucleus accumbens. It is this input that the fMRI signal picks up, as mentioned before. The evidence in Figure 9 also suggests

that the response that Knutson e.a. [2001] recorded in the nucleus accumbens in the short period between stimulus and reward presentations need not be attributed to low time resolution of the fMRI signal. Instead, it is to be viewed as a delayed response to risk (which, in their setting, correlated perfectly with expected reward).

Figure 10 is the analog of Figure 8 for delayed, sustained activation in the nucleus accumbens as a function of risk. The Figure contrasts low risk with high risk. Figure 10 is based on the F-test for the coefficient μ_4 in Eqn. (*). Therefore, the test checks whether there is a *quadratic* relationship in the *sustained* response of the fMRI signal to the conditional variance of reward.

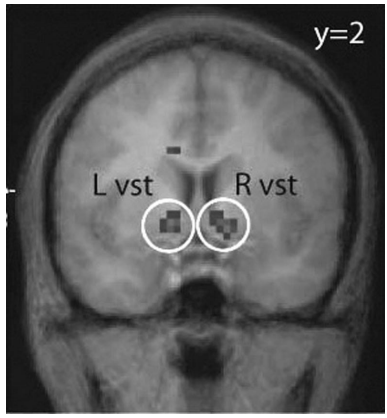

Figure 10. Coronal cross-section of the brain through the nucleus accumbens (ventral striatum; vst) with regions of significantly higher sustained activation for low vs. high conditional variance of winning from 2 seconds after display of the first card until presentation of the second card.
L = left ; R = right. Source: Preuschoff, Bossaerts and Quartz [2006].

IV. Perspective

We have discussed how activations in sub-cortical structures of the human brain correlate with expected reward and risk. As a result, the brain encodes, in at least one region, the parameters that have become the cornerstone of modern finance. It is worth emphasizing at this point that the brain's reaction is almost immediate. It would require far more time for a subject to explicitly compute, for instance, reward variance, than the time it takes these sub-cortical structures to activate. Some of our subjects may never have been exposed to the concept of variance, and hence, may not even be able to consciously compute it!

The locations of the activations that we described here merit emphasis. The crisp signals for expected reward and risk occur in brain regions that are below the cortex. These brain regions are common to many species, including rats and mice, and hence, are often described as conveying "animal spirits," with the idea that irrationality emanates from these regions – while "controlled" actions originate in the (frontal parts of the) cortex (Camerer, e.a. [2005]). This is not to say that the sub-cortical regions cannot be manipulated to produce anomalous behavior such as addiction (Redish [2006]). But our findings demonstrate that the view that the brain is composed of two basic regions, one, primitive and irrational and another one, more recent and rational, is incorrect. In addition, it is not necessarily the primitive part of the brain that causes anomalous behavior. But that should not be surprising: there is a strong evolutionary argument against the conjecture that the primitive part of the brain causes all irrationality.

The same evolutionary argument could also be used to conjecture that emotions are not necessarily detrimental. In the context of reward and risk perception, however, little is known about the value of emotions. It is well known that a financial context triggers emotions (Lo and Repin [2002]). But it remains to be demonstrated that these emotions enhance perception of rewards and risks, or interfere with it. In the area of vision, there are clear examples that emotions enhance perception. See Ling, e.a. [2004]. In the context of decision making, the theory that assigns a positive role to emotions is known as the *somatic marker hypothesis* (Bechara and Damasio, [2005]). Recent research suggests that sound decision making relies on good expression, recording and even simulation of emotions, and that the *amygdala,* the *insula* and *orbitofrontal cortex* all play a crucial role.

Evolutionary arguments may not be relevant for finance, though. Financial risks are too recent to have impacted the human brain. While the brain may have been optimized to evaluate environmental risks, the latter are known to be very different from financial risks. For instance, environmental risks usually do not exhibit the leptokurtosis and (time series) independence that are generally found in financial data. Mistakes are bound to arise if, to evaluate financial risks, the brain invokes processes that are meant to evaluate environmental risks. Human choice of financial prospects indeed exhibits many cognitive biases. It is not known whether these have a physiological basis. If they do, the argument of evolutionary mis-match may explain why.

Note that the activation correlating with risk arrives with a slight delay (we modeled it to be 1 second). It is not known whether there are behavioral implications. If the risk signal in nucleus accumbens is the signal on which choice between gambles is based, then the delay could imply that decisions would be different when taken under severe time constraints (like in fast financial markets, where "good deals" may remain available only for a fraction of a second). It is possible, however, that fast decision making is based on a risk signal elsewhere in the brain. In particular, activation in insula (a part of the cortex) has been shown to also increase with risk, as mentioned in the Introduction. See Huettel, e.a. [2005]. But closer inspection of the data reveals that the insula activation in response to risk is delayed as well (Huettel, e.a. [2005], Figure 3).

Choices between risky gambles should be based on a trade-off between risk and reward. This means that the separate signals for expected reward and risk that we discovered in the brain need to be re-combined somehow to, at a minimum, measure the conflict that exists between these two features of the payoff on a gamble. It is not known how the brain goes about "binding" the two signals. The fact that the activations for expected reward and risk have a different time stamp makes binding obviously more difficult.

This paper discussed one step towards a better understanding of how the brain processes risk and reward. It illustrates that, at least as far as simple gambles are concerned, the brain uses the mathematical representation that has become the foundation of modern finance. That the brain uses mathematics is not a new finding. In the area of vision, for instance, the brain is known to decompose images in terms of, e.g., color, form and motion, compute estimates of these, and integrate these into a coherent perceptual experience that is the brain's best estimate of the physical environment. Not surprisingly, discussions have long centered on how the separate signals are integrated to make sense of the whole (e.g., Singer and Gray [1995]).

Neoclassical and Behavioral Finance treat humans as "black boxes." Our findings illustrate that one can discover surprising things when opening the "black box." Take, for instance, the claim that humans care more about downside than about upside risk (e.g., Ang, Chen and Xing [2006]). This claim is based on behavioral observation. In contrast, variance is a measure of risk that weighs upside risk as much as downside risk. By showing that the brain encodes reward variance, we demonstrated that humans are sensitive to upside risk as much as they are to downside risk – even if they may not always act upon it.

The tools and methodologies we have reviewed here are rapidly evolving, and with them, the collaborations across neuroscience, psychology, economics, and finance. The valuable insights that emerge make it increasingly likely that many of the key questions of finance will be re-visited by looking inside the black box that is the human brain. For instance, our laboratories have recently been investigating changes in perception of risk and reward when participating in asset markets with insiders. The investigation is a truly interdisciplinary endeavor, involving theoretical and experimental finance, neuroscience, systems engineering and philosophy of mind. The results do not fail to provide unique perspectives on how humans attempt to solve extremely complex problems such as extracting information from live market data in the absence of detailed knowledge of the environment.

References

ALLAIS M. [1953], "Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l' école américaine," *Econometrica* 21, p. 503-546.

ANG A., CHEN J., XING Y. [2006], "Downside Risk," *Review of Financial Studies*, forthcoming.

BECHARA A., DAMASIO A. D. [2005], "The somatic marker hypothesis: A neural theory of economic decision," *Games and Economic Behavior* 52, p. 336-372.

BOSSAERTS P., PLOTT C. [2004], "Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets," *Review of Finance* 8, p. 135-169.

BULLOCK T. H., BENNETT M. V., JOHNSTON D., JOSEPHSON R., MARDER E., FIELDS R. D. [2005], "Neuroscience. The Neuron Doctrine," *Science* 310, p. 791-793.

CAMERER C., LOEWENSTEIN G., PRELEC D. [2005], "Neuroeconomics: How Neuroscience Can Inform Economics," *Journal of Economic Literature* 43, p. 9-64.

CRITCHLEY H. D., MATHIAS C. J., DOLAN R. J. [2001], "Neural Activity In The Human Brain Relating To Uncertainty And Arousal During Anticipation," *Neuron* 29, p. 537-545.

DAYAN P., YU A. J. [2003], "Uncertainty And Learning," *IETE Journal of Research* 49 (2), p. 1-10.

FIORILLO C. D., TOBLER P. N., SCHULTZ W. [2003], "Discrete Coding of Reward Probability and Uncertainty by Dopamine Neurons," *Science* 299, p. 1898-1902.

FIORILLO C. D., TOBLER P. N., W. SCHULTZ [2005], "Evidence that the delayperiod activity of dopamine neurons corresponds to reward uncertainty rather than backpropagating TD errors," *Behavioral and Brain Functions* 1:7 (15 June 2005).

HSU M., BHATT M., ADOLPHS R., TRANEL D., CAMERER C. F. [2005], "Neural systems responding to degrees of uncertainty in human decision making," *Science* 310, p. 1680-1683.

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, p. 3304-3311.

HUETTEL S. A., STOWE, C. J. GORDON E. M., WARNER B. T., PLATT M. L. [2006], "Neural Signatures of Economic Preferences for Risk and Ambiguity", *Neuron* 49, p. 765-775.

KNUTSON B. K., FONG G. W., BENNETT S. M., ADAMS C. M., HOMMER D. [2003], "A Region Of Mesial Prefrontal Cortex Tracks Monetarily Rewarding Outcomes: Characterization With Rapid Event-Related fMRI", *NeuroImage* 18, p. 263-272.

LAURY S., HOLT C. [2002], "Risk Aversion and Incentive Effects", *American Economic Review* 92, p. 1644-1655.

LING S., PHELPS E., HOLMES B., CARRASCO M. [2004], "Transient Attention Increases Perceived Saturation", *Vision Sciences Society*, Sarasota, Fl. (poster).

LO A., REPIN D. V. [2002], "The Psychophysiology of Real-Time Financial Risk Processing," *Journal of Cognitive Neuroscience* 14, p. 323-339.

LOGOTHETIS N. K., PAULS J., AUGATH M., TRINATH T., OELTERMANN A. [2001], "Neurophysiological Investigation Of The Basis Of The fMRI Signal," *Nature* 412, 12 July 2001, p. 150-157.

MARKOWITZ H. [1952], "Portfolio selection," Journal of Finance 7, p. 77-91.

MONTAGUE P.R., DAYAN P., SEJNOWSKI T.J. [1996], "A Framework For Mesencephalic Dopamine Systems Based On Predictive Hebbian Learning," *Journal of Neuroscience* 16, p. 1936-1947.

PLATT M. L., GLIMCHER P. W. [1999], "Neural Correlates Of Decision Variables In Parietal Cortex", *Nature* 400, 15 July 1999, p. 233-238.

POHL R., Ed. [2005], *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*, New York: Psychology Press.

PREUSCHOFF K., BOSSAERTS P., QUARTZ S. [2006], "Neural Differentiation of Expected Reward and Risk in Human Subcortical Structures," *Neuron,* in press.

REDISH A. D. [2004], "Addiction as a Computational Process Gone Awry," *Science* 306, 10 December 2004, p. 1944-7.

TANAKA S., DOYA K., OKADA G., UEDA K., OKAMOTO Y., YAMAWAKI S. [2004], "Different Cortico-Basal Ganglia Loops Specialize in Reward Prediction on Different Time Scales," working paper.

SEYMOUR B., O'DOHERTY J., DAYAN P., KOLTZENBURG M., JONES A. K., DOLAN R. J., FRISTON K. J., FRACKOWIAK R. S. [2004], "Temporal Difference Models Describe Higher-Order Learning In Humans" *Nature* 429, 10 June 2004, p. 664-667.

SCHULTZ W., DAYAN P., MONTAGUE P. R. [1997], "A Neural Substrate Of Prediction And Reward," *Science* 275, p. 1593-1599.

SINGER W. and GRAY C. M. [1995], "Visual Feature Integration and The Temporal Correlation Hypothesis," *Annual Reviews of Neuroscience* 18, p. 555-586.

TOBLER P. N., FIORILLO C. D., SCHULTZ W. [2005], "Adaptive Coding of Reward Value by Dopamine Neurons", *Science* 307, 11 March 2005, p. 1642-1645.

VOLZ K. G., SCHUBOTZ R. I., VON CRAMON D. Y. [2003], "Predicting Events Of Varying Probability: Uncertainty Investigated By fMRI," *NeuroImage* 19, p. 271-280.