

Information processing and decision-making: evidence from the brain sciences and implications for Economics*

Isabelle Brocas

University of Southern California and CEPR

March 2012

Abstract

This article assesses the potential benefits of including findings from neurobiology in economic decision-making models. First, we emphasize that the evidence supports both ‘expected utility-like’ theory and ‘Bayesian-like’ information acquisition theory. Second, we explain that inferences and representations are subject to physiological limitations that affect decision making. We report in particular two ‘mechanical’ models developed in neuroscience to represent neural data and choices. We then propose two economic models that incorporate physiological limitations into an expected utility framework. Interestingly, these two models provide foundations for those developed in neuroscience (which emerge endogenously in our framework) and provide further predictions that can be tested in principle. This allows us to discuss the benefits of bringing together evidence from neuroscience and economic modeling.

*Correspondence address: Department of Economics, University of Southern California, 3620 S. Vermont Ave. Los Angeles, CA 90089-0253, USA, email: <brocas@usc.edu>. I am grateful to Juan D. Carrillo and two anonymous referees for very useful comments. I thank Bo Ren for research assistance.

1 Introduction

Economic models of decision-making rely on several important paradigms. Two of them are particularly important. First, decision-makers' choices can be represented by well-defined utility functions (satisfying well-defined axioms). Second, decision-makers are Bayesian processors of information. Even though these two paradigms have merits and surely account for a large proportion of observed behaviors still, various anomalies have been reported. Most of the hypotheses formulated to justify those anomalies are introspective and make implicit claims about the underlying mechanisms leading to behavior. In parallel, a myriad of experiments in neuroscience try to describe the brain processes involved in decision-making. We believe that this evidence could and should shed light on the fundamental causes of observed choice anomalies. The objective of this article is to provide a motivated summary of a few findings from neurobiology and describe how they can be incorporated in economic models in a useful way.

We are interested in individual decision-making in the following simple experiment. There are two possible states of nature $S \in \{A, B\}$. A decision-maker (hereafter DM, he) faces uncertainty regarding the state. Nature chooses state A with probability p_0 and state B with probability $1 - p_0$ but the state is not communicated to DM, who must decide between two possible actions $k \in \{a, b\}$. The benefit of taking action a is G_A if the true state is A and 0 if the true state is B . Similarly, the benefit of taking action b is G_B if the true state is B and 0 if the true state is A . DM may obtain some information, a signal s correlated with the true state S , prior to making his decision.

From the perspective of economists, this is a simple problem of information processing: DM receives a signal s , he then uses this information to form an updated belief about S , and makes a decision based on this belief. Economists would further assume that the updating process is Bayesian and the decision is made by maximizing expected utility. They also would model the signal as an imperfect indicator of the true state, and assume that the information transmitted is equal to the information processed.

This paradigm follows closely standard experimental designs in Neurobiology. A typical experiment is the well-known random dot stimuli experiment (see for example Newsome et al. (1989) and Salzman et al. (1990)). A monkey (DM) is facing a black computer background showing moving white dots. The percentage of dots that move in the same direction, or *coherence*, can be manipulated by the experimenter. A 100 % coherence corresponds to the case they all move at the same speed and direction. A 0% coherence corresponds to the case they all move randomly. The more or less coherent dots display corresponds to s in the previous notations, and it contains information about the true

direction, which corresponds to S . The monkey must identify the direction followed by the majority of dots by making a saccade. Correct answers are rewarded.

Neurobiologists work under three assumptions. (i) The stimulus is encoded by neurons in the *sensory system* (e.g. visual and auditory systems), that is, a representation of s is created. This representation is data available to the brain, which we call D . (ii) Information is then interpreted to determine which response should be triggered. (iii) The decision is implemented by the *motor system*. However, neurobiologists do not presuppose how information is processed or how decisions are made. Rather, they measure cell activity in preselected brain areas and correlate it with response.

In this article, we use evidence from neurobiology to understand how DM processes the information contained in the signal and makes a decision. This allows us to revisit existing economic models of decision-making and offer representations more in line with the evidence. For expositional purposes, we will use the experimental two-action discrimination paradigm described above.¹ We will also restrict attention to single neuron studies in non-human primates. This decision is somewhat arbitrary and the reader shall keep in mind that many other kinds of experiments have been conducted, including single neuron studies on rodents, fMRI experiments on humans and lesion studies on humans, among others.

The article is organized as follows. In section 2, we describe the underlying mechanics of information transmission across brain cells. In section 3, we investigate the implications of the existing evidence for expected utility and Bayesian learning. In section 4, we review decision-making models developed in neuroscience. In section 5, we present economic decision-making models inspired by the evidence reviewed in sections 2 and 3 and the models outlined in 4. Section 6 concludes.

2 The underlying mechanism

For the reader unfamiliar with brain functioning, it might be useful to start with a few words on anatomy and physiology. This description is by no means exhaustive and should be completed by further readings.²

2.1 Decision-making and brain systems

The human brain has a wide array of functions including functions to help us interpret sensory inputs, originate and coordinate motor responses, control basic functions such as

¹Most of the literature focuses on such two-action discrimination tasks.

²See for instance Thompson (2000), Shepherd (2003) and Nicholls et al. (2001).

sleep, hunger or thirst, stock and retrieve information required to perform tasks, and guide abstract and complex decision-making. The cerebral cortex controls higher functions. It is composed of various systems with specific missions.

The primary sensory cortices are involved in the treatment of sensory information.³ Representations are created from primary sensory inputs, which corresponds to the transformation of s into D in our experiment.

The motor cortices (primary and secondary) are responsible for the planning and execution of voluntary motor functions. They host the process of action implementation rendering decisions. This corresponds to implementing action a or action b .

The prefrontal cortex is in charge of complex cognitive tasks, expressing personality, and moderating behavior.⁴ The primary activity of the prefrontal cortex is thought to be cognitive control, that is the orchestration of thoughts and decisions in accordance with internally generated goals.⁵ An additional set of structures, sometimes called the limbic system,⁶ support a variety of functions including the regulation and manifestation of emotions and the formation of memory.⁷ These structures participate in the formulation of a decision for a typical problem.

Neurophysiological studies reveal that the sensory and motor systems of the cerebral cortex are crucial structures (see Sugrue et al. (2005) for a review) in tasks in which the monkey is trained to indicate a decision with an eye movement. The transformation of s into D is made in areas of the occipital cortex where sensory neurons encode visual stimuli. Decision-making is also made in areas of the parietal and frontal cortices. Decisions are based on the sensory stimuli strength and decision-related systems make the animal's ultimate choice.

³These include the primary visual cortex (in the occipital lobe), the primary somatosensory cortex (in the parietal lobe), the primary auditory cortex (in the temporal lobe), the primary olfactory cortex (in the piriform cortex in the temporal lobe) and the primary gustatory cortex (in the anterior insula and the frontal operculum). Higher order sensory processing occurs also in the orbitofrontal cortex. I refer the reader to Thompson (2000) for a more detailed overview.

⁴It is composed of the orbitofrontal (OFC), the ventromedial (vm-PFC), the dorsolateral (dl-PFC), the ventrolateral (vl-PFC), the medial (m-PFC), the anterior (a-PFC) as well as the cingulate cortex.

⁵See the theories of cognitive control and, in particular Miller and Cohen (2001).

⁶The boundaries of the limbic system are difficult to delineate and there has been controversy over the terminology. More recently, it has been suggested that the concept should be considered obsolete. See Phelps (2009).

⁷These are the hippocampus, responsible for the formation of long term memory; the amygdala, performing an important role in the processing and memory of emotional reactions; the hypothalamus, responsible for controlling body temperature, hunger, thirst, fatigue, anger, circadian cycles and maintaining overall homeostatic balance; and the thalamus, involved in the regulation of sleep and a primary relay station to translate sensory information in a form readable for the cerebral cortex.

2.2 Neural networks and information transmission

The brain is composed of many different types of cells but the main functional unit is the neuron.⁸ It is organized in a network involving around ten billion neurons. Neurons within the brain process and transmit signals. Some convey information to the nervous system while others transmit signals from the nervous system. Even though neurons are highly specialized and there is a wide variety of them, a typical neuron has a central part (the soma) and cellular extensions (dendrites) that form a tree. Dendritic trees are the place where the neuron receives information. The signal then passes through the soma and travels away from it down the axon, which carries nerve signals to another cell. The axon may divide into a large number of branches enabling communication with many target neurons. The terminal of the axon contains synapses that are either electrical or chemical. Electrical synapses are direct and electrically conductive junctions between cells. Chemical synapses release neurotransmitter chemicals to communicate with target neurons.⁹

If a neuron responds to a nerve impulse, then it responds completely (all-or-none principle). Several events must occur for a neuron to respond and pass the signal to other cells. When a stimulus is received (after neurotransmitter release by a presynaptic neuron), the membrane is depolarized, that is its voltage is increased. This triggers the opening of sodium and potassium ion channels located in the axon. For small voltage increases, the potassium current dominates and the voltage returns to its normal resting value.¹⁰ However, if the voltage increases past a critical threshold higher than the resting value, the sodium current dominates.¹¹ The positive feedback activates even more sodium channels¹² and the cell “fires”, producing an electrical impulse, called action potential or spike, that travels down the axon.¹³ It arrives at the synapse and produces an influx of

⁸Neurons are outnumbered by glial cells which are now also receiving attention in the literature.

⁹There exist many different neurotransmitters. Acetylcholine is usually an excitatory neurotransmitter and governs muscle contractions and causes glands to secrete hormones. Gamma-aminobutyric acid (GABA) is usually an inhibitory neurotransmitter and can be found in the motor and visual systems. Dopamine can have an excitatory or inhibitory action on the post-synaptic cell, and its effect depends on the receptor present on that cell. It is involved in mood and the control of complex movements. Dopamine-sensitive neurons are located in the PFC. Other important neurotransmitters are serotonin, glutamate, adenosine and adenosine triphosphate (ATP).

¹⁰See Hodgkin and Huxley (1952) for a seminal model of voltages and currents of the action potential.

¹¹The critical threshold voltage is usually around -45 mV, but it depends on the recent activity of the axon. A membrane that has just fired an action potential cannot fire another one immediately, since the ion channels have not returned to their usual state. The period during which no new action potential can be fired is called the absolute refractory period. After some but not all of the ion channels have recovered, the axon can be stimulated to produce another action potential, but only with a much stronger depolarization, e.g., -30 mV. Thresholds can be modified by secondary messengers (like hormones or neurotransmitters) that also communicate with the neuron.

¹²This “summation” of incoming information is made at the axon hillock.

¹³Some neurons do not generate action potential, but instead they generate a graded electrical signal.

calcium ions that trigger a biochemical cascade resulting in neurotransmitter release.

Neurons can be classified according to their discharge patterns. Some are constantly active (e.g. some interneurons). They produce a steady rate of firing and respond to increase intensity of a stimulus by increasing their firing frequency. Other neurons are phasic and firing decreases or stops when the stimulus becomes steady. Some neurons have fast firing rates (fast spiking neurons) while others have narrow action potentials (e.g. thin-spike interneurons in the prefrontal cortex). Synapses can be excitatory, inhibitory or modulatory, depending on which effect they have on their target neurons.¹⁴ The post-synaptic receptor is responsible for the action of the neurotransmitter released at the synapse. In other words, synapses will increase or decrease the activity in the target neuron. When excitatory synapses are activated, the target neuron is induced to produce nerve impulse. By contrast, when inhibitory synapses are activated, the action of excitatory synapses is counterbalanced to prevent nerve impulse. The total activation of synapses determines when nerve impulses are sent out to other neurons. The strengths of synapses (i.e. the effectiveness with which synapses produce or negate nerve impulse) vary largely. Information sent via strong (resp. weak) synapses is emphasized (resp. de-emphasized). Many neuronal mechanisms are in place to modify strength. Some cause persistent alterations that can last months or years, others cause short-term modifications lasting from fractions of seconds to a few minutes.¹⁵

Measuring cell activity informs researchers about the involvement of brain areas into specific tasks. Studies rely on such measures and draw correlations between cell activity and decision-making. In single neuron studies, two types of measures can be found: spike trains and firing rates. The fundamental unit of information in the brain is the action potential. However, given the all-or-none principle, an action potential per se does not carry much information. Rather, the timing of action potentials does. For instance, increasing the intensity of the stimulation does not produce a stronger signal but can produce more impulses per second. The response of a single neuron can be summarized as a list of dates at which the neuron fired action potentials, or spike train. Alternatively, we can compute the average number of action potentials per unit of time, the firing rate.¹⁶

They are called non spiking neurons. They are generally sensory neurons or interneurons.

¹⁴For a thorough description of synaptic functioning, see Cowan et al. (2003). Note that it is widely believed that synaptic strength provides the basic mechanism for learning and memory. See Hebb (1949) for a seminal model of learning and memory.

¹⁵These mechanisms are referred to as short-term plasticity. Short-term plasticity is correlated with the history of the activity of a synapse. It provides a way for a neuron to emphasize or de-emphasize particular types of information according to the pattern of activation in the past. It is a history dependent filter that decides which information is the most important in determining the output of a neuron. See for instance Klyachko and Stevens (2006) for a study of hippocampal synapses.

¹⁶This measure implicitly assumes that the exact timing of spikes should play no role. For a discussion,

2.3 Summary

Decision-making involves brain systems that process the information made available to DM towards a decision that is eventually implemented. Systems are composed of neurons tuned to react to specific characteristics of the information they receive from the outside world or from neurons in other systems. Neuronal cell-firing and synaptic connectivity determine how a sensory stimulus (received in the sensory cortices) will be interpreted and mapped into a response (implemented by the motor cortex). Throughout the process, some information will be emphasized or deemphasized, as synapses act as a filter of information.

3 Representations

In this section, we relate findings in neurobiology to the two building blocks of Economic Theory described in the Introduction, namely Expected Utility and Bayesian updating. We want to assess whether these two paradigms are reliable representations of DM's motivations and ability to reason.

From the perspective of Economics, a DM facing our experiment should evaluate action a at $P(A)G_A$, where $P(A)$ is the probability that the true state is A . Similarly, he should evaluate action b at $P(B)G_B = (1 - P(A))G_B$. It is therefore efficient to choose action a (respectively b) if $P(A)G_A > (1 - P(A))G_B$ (respectively $P(A)G_A < (1 - P(A))G_B$). Now, to formulate a decision, the brain must represent and combine two pieces of information. First, information D needs to be extracted from the signal s to represent probabilities $P(A)$ and $P(B)$. Second, the magnitude of the payoffs G_A and G_B must be represented.

3.1 Expected utility

Expected utility theory presupposes that DM is able to represent uncertainty as well as the magnitude of rewards, and evaluate actions in expectations. To test this hypothesis, Platt and Glimcher (1999) designed an experiment in which information about the magnitude of a reward as well as the probability of the reward are displayed visually to rhesus monkeys. Subjects are then asked to choose an alternative and are rewarded accordingly.

This exercise reduces our experimental paradigm to a task that relaxes the signal extraction problem. More precisely, DM is "given" the probabilities of the states $P(A)$ and $P(B)$ (and does not need to infer them from noisy signals) and must make the correct decision as a function of the magnitude of these probabilities and the magnitude of the

see Abbot (1994) and Gerstner and Kistler (2002).

rewards G_A and G_B . This design allows to test directly for expected utility theory by manipulating probabilities and rewards.¹⁷

Single intraparietal neurons were recorded, as they are sensitive to the probability that a particular response will result in a gain. The authors showed that rewards as well as the probability that a particular choice result in a reward modulate the activity of the neurons in the lateral intraparietal area. Moreover, when subjects are free to choose between two responses, the expected gain influenced both the choice and the activation of parietal neurons. In other words, probabilities and rewards were represented. Both activation and choice were also roughly compatible with expected payoff maximization.¹⁸ That is, cell activation and behavior are in line with the expected utility representation. Said differently, DM is capable of computing and comparing the expected payoffs of actions a and b .

A myriad of subsequent studies have reported similar findings and tried to uncover the circuitry involved in value representation and action selection. Among others, Padoa-Schioppa and Assad (2006, 2008) studied the representation of economic value by analyzing the activity of neurons in the orbitofrontal cortex of the monkey. The authors found that those neurons encode the value of goods that are offered and chosen. Shortly after the offer, neurons encoding the value of one alternative or the other are active. Then, neurons encoding the value of the chosen alternative are active, independently of the characteristics of this alternative. Last and after the choice is revealed, neurons encode the identity of the chosen alternative.¹⁹ Sugrue et al (2004) studied the subjective evaluation of alternatives by measuring the activity of neurons in the parietal cortex. Relying on a dynamic foraging task in which the values of alternatives changed over time, they found that the relative value of the alternatives were represented by these neurons. Samejima et al. (2005) found that action values were represented in the striatum (basal ganglia). Lau and Glimcher (2008) reported that some neurons were responsible for encoding chosen values. More precisely, the authors found in the striatum neurons encoding action values and neurons encoding chosen values.

Taken together, these studies show that the brain is able to represent both ‘objective’ rewards offered by the experimenters and ‘subjective’ values such as tastes. Neurons in the parietal cortex, the prefrontal cortex and the basal ganglia (among others) are part of the network linking visual stimulation to oculomotor behavior, and encode the

¹⁷Two DM may represent the rewards in different ways, e.g. if they have different risk attitudes.

¹⁸See also Glimcher and Rustichini (2004) for a review of studies offering insights on how (and where) the desirability of options is represented.

¹⁹Moreover, the activity of a neuron in response to an alternative seems to be independent of the menu of alternatives that is offered.

expected values of alternatives (see McCoy and Platt (2005) as well as Kable and Glimcher (2009) for reviews). Several studies have shown that other areas also encode reward-related information. For instance, the recent literature on reward-prediction (see Schultz (2006)) shows that midbrain dopaminergic neurons signal the difference between reward expectation and reward outcome in learning environments. For our simple experiment, the evidence just reviewed validates expected utility theory. We gather this important finding in the next result.

Result 1 *For given $P(A)$, G_A and G_B set by the experimenter, values and probabilities are represented (in the orbitofrontal cortex and the striatum) and alternatives are compared (in the lateral intraparietal area). Decision-making is consistent with choosing the action that maximizes DM’s expected utility. Formally, the optimal action k^* satisfies*

$$k^* = \begin{cases} a & \text{if } P(A)G_A > (1 - P(A))G_B \\ b & \text{if } P(A)G_A < (1 - P(A))G_B \end{cases}$$

3.2 Inferences

Random dot stimuli type of experiments are meant to study how DM extracts information from a noisy signal. In such experiments, the probability of the reward must be assessed by DM. The experimenter manipulates the signal s sent to DM but typically not the magnitude of the reward. The purpose of this section is to understand how information about s is processed in the primary sensory cortices and used to produce a representation of the likelihood of the events. Compared with studies in the previous section, the experiments now assume $G_A = G_B$ and focus on the information that DM infers from s .

There exists a fundamental difference between encoding and decoding. Encoding corresponds to the transformation of s into neural responses. Decoding consists in estimating the true state S using neural responses. The overall process is implemented by different layers of neurons. Neurons in the sensory cortices encode information, which produces a response (measured in spike trains) D . The activity can be described by the conditional probability distribution $P(D|S)$. The next layer of neurons extracts this information by ‘calculating’ an estimate of S based on D . It is important to note that decoding requires to measure the deviations between the estimated stimuli and the actual stimuli. It is therefore necessary to assess how the activity of neurons in the first layer correlate and how these correlations are related to the stimulus. We review some findings in the next two sections. We refer the interested reader to Schneidman et al. (2003) for details about the difficulties arising in decoding noisy information.

3.2.1 Encoding

The signal sent by the experimenter is received by the sensory system where it is encoded. Given a stimulus s reflecting S , the response of neurons is given in terms of a probability distribution $P(D|S)$. For instance, in the random dot stimuli experiment, S would be the true direction of dots and s the more or less coherent display. In that setting, D represents neural activity in the area of neurons that encode s . Researchers take repeated measurements of the response of the neuron, which generates a curve describing the responses (Figure 1). This curve simply represents the likelihood of obtaining each possible response $P(D|S)$. We can derive such a curve for each hypothesis, which provides us with two probability measures $p(D|A)$ and $p(D|B)$.²⁰

[INSERT FIGURE 1]

Moreover, each neuron has a preferred direction (e.g. some are best at detecting upwards movements, others are best at detecting downwards movements). These preferred directions correspond to the possible hypotheses S . In Figure 2, the x -axis represents the preferred direction of neurons (A and B) and the y -axis represents the number of spikes. Each dot corresponds to the response of a neuron: the x coordinate is its preferred direction and the y coordinate the number of times it spiked after receiving s .

[INSERT FIGURES 2A AND 2B]

Many analyses have shown that the neural response D is variable. Variability reflects the fact that the data used by the brain is not fully accurate and in particular that neurons do not fire in a fully consistent way. This can be seen in Figures 1 and 2. For a given stimulus, the response of a neuron is variable and neurons with the same preferred directions do not all show the same activity. A few hypotheses have been formulated leading to a myriad of experiments and theories. In particular, studies have shown that neurotransmitter release is stochastic and synaptic release probabilities are heterogeneous across synapses, sometimes dependent on the history of the activity (see Stevens (2003) for a review). Also, the reliability of the data depends on experimental conditions. For instance, firing distributions may differ in naturalistic conditions compared to other conditions (see Simoncelli (2003), Lewen et al. (2001)). Part of the noise in neural signals

²⁰Several theories based on information theory (Shannon (1948)) have been proposed to explain the goal of early stages of visual processing. According to the Efficient Coding Hypothesis (Barlow (1961),(2001)), it is to produce an efficient representation of the incoming visual signal. His model treats the sensory system as a communication channel where neuronal spiking is an efficient code to represent s . To be efficient, the code must maximize the channel capacity and minimize the redundancy in the representation. See also the predictive coding theory (Srinivasan et al. (1982)), the maximal decorrelation theory (Atick and Redlich (1992)), the maximal signal-to-noise ratio theory (McCarthy and Owen (1996)), and the maximal kurtosis theory (Field (1994)).

is due to the imprecision around the stimulus s and the context in which s is displayed affects the perception about the true state of the world S .

Result 2 *Information encoded in the primary visual cortex is subject to noise. This variability affects the distribution of the neural responses conditional on the true state $P(D|S)$.*

For our experiment, the main lesson is that internal and external noises act on D as there is an imperfect correlation between s and S as well as an imperfect correlation between D and s .²¹

3.2.2 Representation of prior probabilities

In a recent study, Hanks et al. (2011) showed that information related to prior probabilities is combined with sensory evidence. To do so, the authors used a random dot stimuli task and they manipulated the prior probabilities of rewarded choices. They found in particular that prior probabilities were incorporated into the decision process and was reflected in the firing rate of neurons in the lateral intraparietal area. This result is consistent with Platt and Glimcher (1999).

3.2.3 Aggregation of individual neuronal responses and inferences

Sensory activity is encoded by the activity of a large population of neurons, and individual neuronal responses need to be aggregated into a representation. For instance, neurons performing early detection produce an activity consistent with the graph in Figure 2b. This information is then sent to a new layer of neurons. The purpose of this section is to analyze the way these neurons interpret the information they receive.

Stimuli are processed in different areas. For instance direction stimuli are processed in area V1 (striate cortex) and area MT (middle temporal visual area). MT neurons are specialized in motion processing and are not sensitive to a broad range of stimuli like color or form. Their neuronal responses to direction are generally stronger than V1 neurons (see Albright (1984)). This suggests that some features may be emphasized in some areas and at the same time de-emphasized in others. Also, some areas receive projections from other areas that have already performed a transformation of the incoming information. For instance, area MT receives projections from V1 and V2 areas, along with several others (see Van Essen et al. (1992)).

²¹This latter noise is often assumed to be Poisson, because the variance increases with the mean activity, a feature consistent with the activity of cortical neurons.

It has been shown that the prevalent reaction might result from either competition or cooperation between neurons with different perspectives. For instance Nichols and Newsome (2002) use a random dot stimuli task to study neurons in the visual area with different preferred directions (e.g. detect accurately their preferred direction and less well other directions). The authors show that neurons that have a wide range of preferred directions can cooperate to shape estimates for directions (consistent with a vector averaging model), while those with opposite preferred directions compete (consistent with a winner-take-all model). A classical model is the population vector estimator. It assigns to each neuron a vector. Its length is proportional to the neuron’s activity and its direction corresponds to its preferred orientation. Then, it sums all the individual vectors to form a population vector, and estimates the orientation from the angle of this vector (see Georgopoulos et al. (1986) for an application to neuronal activity in the motor cortex and Gielen et al. (1988) for the case of the auditory nerve). In the example of Figure 2b, this would result in fitting the neuronal activity with a cosine function, as it is represented by the solid curve in Figure 3a. The peak of the function is the estimate of the true state \hat{S} . Other approaches have been proposed to represent neuronal activity of large populations of neurons. One common method is the maximum likelihood (ML) estimator (see Abbott (1994) for a review). Neural firing rates for a stimulus s are treated as random variables, and an overall estimate (the likelihood) of the probability of the set of observations is taken. This method results in producing the solid curve in Figure 3b. The peak is the ML estimator \hat{S} .²² Deneve et al. (1999) show that biologically plausible networks can implement ML.²³

[INSERT FIGURE 3A AND 3B]

According to Ma et al. (2006), the response of neurons is given in terms of the probability distribution $P(D|S)$ that encodes the posterior distribution $P(S|D)$ over the true state S . The Bayesian posterior belief of state A following the observation of D is

$$P(A|D) = \frac{p(D|A)}{p(D)} p_0 \propto P(D|A) p_0$$

and therefore the probability distributions can be recovered from neuronal responses.²⁴ Ma et al. (2006) also show that Bayesian inferences are simple to implement by population codes with Poisson-like variability. Indeed, if variability is Poisson, simple linear

²²The ML estimator is found to perform better as it has a minimum variance over trials under some conditions.

²³In other words, if the network receives noisy pooled input from a population of neurons, the output activity of the network is an optimal ML estimator.

²⁴The prior belief p_0 could be encoded by a population of neurons that spike before the beginning of the trial.

combinations of populations of neural activity (say a population encoding visual cues and a population encoding auditory cues) achieve Bayesian inferences.²⁵

Result 3 *The brain is capable of drawing approximately Bayesian inferences from signal. More precisely, the activity of neurons D is given as a probability distribution $P(D|S)$ that encodes posterior beliefs about the hypotheses $P(S|D)$.*

3.3 Summary

Result 1 shows that given a set of beliefs about the true state of the world and related rewards, the brain can compute the expected payoff of each alternative and compare them. Result 3 establishes that the brain is capable of estimating those beliefs accurately. In particular, it can extract information from a signal and encode Bayesian posterior beliefs about the true states of the world.

Formally, a DM facing our experiment will evaluate action a at $P(A)G_A$, where $P(A)$ is the posterior probability that the true state is A conditional on the signal s sent by the experimenter and the ‘data’ D encoded. In previous notations, $P(A) = P(A|D)$ where D is a noisy representation.

In the case of the random dot stimuli task, a very coherent display s where the majority of dots move in the same direction (state A) will generate high activity from neurons tuned to detect that direction and a higher probability that DM selects that direction (action a) as a response, both consistent with a high $P(A|D)$. This informs us that DM is confident enough that the dots move actually in that direction to indicate so with a saccade. This also suggests that enough evidence of A triggers action a and insufficient evidence of A triggers action b . In the next section, we review models that refine this idea and address decision-making.

4 Decision-making

So far, we have reviewed evidence that validates (i) the expected utility approach to evaluate alternatives for given representations of probabilities and rewards, and (ii) the Bayesian model to assess the likelihood of S given a noisy representation D of a noisy signal s . In this section, we report two families of models proposed in the neuroscience literature to explain how decisions relate to D , and how DM chooses between alternatives.

²⁵Suppose there are two cues 1 and 2 encoded by two different populations representing likelihood functions $p(D_1|S)$ and $p(D_2|S)$, Bayesian inferences can be achieved by constructing a third population with $D_3 = D_1 + D_2$ as long as the prior is flat and variability is Poisson. This results extends under weaker conditions.

These models do not rely on expected utility theory and optimization. We shall study in section 5 their relationship to it.

4.1 Comparison of signals

In studies based on the random dot stimuli task, researchers compare two sets of data: behavioral data and neural activity. On the one hand, the proportion of correct answers can be plotted for each coherence level. This data can be fitted using standard statistical techniques to obtain a “psychometric function” (Figure 4). This function is a representation of *behavior*. On the other hand, it is possible to observe the number of spikes for each possible coherence level and each trial. Assuming that a decision is made in favor of the decision preferred by the neuron under study if its response is large enough,²⁶ it is possible to estimate the probability that the decision rule give a correct response for each coherence level. This allows to compute a “neurometric function” comparable to the psychometric function. Those functions are usually not significantly different in this type of experiments. The decision favoring the preferred direction of the neuron under study (resp. null direction) predominates when the neuron fires strongly (resp. weakly).

[INSERT FIGURE 4]

Inspired by these results, Shadlen and al. (1996) propose a simple flexible pooling model that links (pooled) neuronal responses to decisions. Formally, suppose there are $2N$ neurons each represented by a random variable, N of which favoring direction A and the remaining N favoring direction B . The response of each group of neurons (spike counts) is then averaged into a signal measure and the larger measure dictates the decision.²⁷ In other words, when the activity in favor of one alternative is higher, then this alternative is chosen. In the same lines, Ditterich et al. (2003) show that, to decide between two alternatives, the response of neurons with opposing preferences are compared to obtain net evidence for one alternative. The sign of the difference represents which alternative is more likely, and the magnitude represents the persuasiveness of the evidence.

This class of models can be roughly summarized as follows. Evidence is encoded into D which might be the list of spike counts of neurons firing in response to s . Decision-making consists in simplifying this information into a measure Δ and act upon this measure. In a two-action discrimination task, some realizations of Δ indicate confidence in A and the others indicate confidence in B . To simplify the exposition, let us introduce an example

²⁶Britten et al. (1992) compare the signal coming from the neuron under study and the signal from an hypothetical anti-neuron with preferred opposite direction.

²⁷The model also assumes the brain cannot compute the average with precision and considers a pooling noise. Pooling has the effect to depreciate sensitivity.

for reference. We assume that A refers to danger and B absence of danger, and we suppose that the relevant information is captured by the activity of neurons in a specific region (e.g. the amygdala). In that example, Δ is a single variable, low activation indicates confidence in B while high activation indicates confidence in A . Therefore, there exists a value threshold Δ^* to which Δ can be compared.

Result 4 *Decision-making is consistent with a threshold model in which evidence is pooled into a measure Δ that is compared to a decision threshold Δ^* . Then $\Delta > \Delta^*$ is evidence of A and a is triggered, while $\Delta < \Delta^*$ is evidence of B and b is triggered.*

Gold and Shadlen (2001) propose a related model of neural computations in which information over time is accumulated into a single quantity: the logarithm of the likelihood ratio (LR) favouring one alternative over the other. As noted earlier (see Figure 1), the response of a sensory neuron (spike rate per second) is a random variable and the neural response varies as a function of the experimental conditions. Given the true state S , the activity D can be summarized in the probability density function $P(D|S)$ similar to the one obtained in Figure 1. In this experiment, the theoretical posterior belief of state A is

$$P(A|D) = \frac{P(D|A)p_0}{P(D|A)p_0 + P(D|B)(1 - p_0)}$$

where $P(A)$ and $P(B)$ are prior beliefs. Interpreting the evidence in favor of A should result when $P(A|D) > P(B|D)$, that is when the likelihood ratio $\frac{P(D|A)}{P(D|B)}$ is high enough:

$$P(A|D) > P(B|D) \Leftrightarrow \frac{P(D|A)}{P(D|B)} > \frac{1 - p_0}{p_0} \quad (1)$$

Suppose now there are two neurons (or populations of neurons): the first detects signals in favor of A (and is more likely to fire if A is the true state), while the other detects signals in favor of B (and is more likely to fire if B is the true state). The response of the first is denoted by D_A and the response of the second is denoted by D_B . In that case the likelihood ratio is:

$$\frac{P(D_A, D_B|A)}{P(D_A, D_B|B)} = \frac{P(D_A|A)}{P(D_A|B)} \cdot \frac{P(D_B|A)}{P(D_B|B)}$$

taking the logarithm of the likelihood ratio linearizes the problem that can now be solved by taking simple additions. Gold and Shadlen (2001) conjectured that LIP neurons accumulate sensory responses from pools of neurons to form a decision variable that approximates log LR. It is sufficient to compare log LR to a threshold that triggers the action consistent with hypothesis A if log LR is above that threshold. Overall, according to Gold and Shadlen (2001), decision-making is consistent with a threshold model in which

evidence is pooled into the log-likelihood ratio favoring A over B , namely $\Delta = \log \frac{P(D|A)}{P(D|B)}$ and compared to a decision threshold Δ^* .

Of course, the threshold Δ^* affects how the evidence is interpreted. Gold and Shadlen (2001) conjecture that the threshold must make sure that the best rewards are obtained as often as possible, and therefore be an indicator of the expected values of the possible responses. More formally, they propose that the left hand side of equation 1 should reflect not only prior probabilities but also accommodate reward values G_A and G_B .

4.2 Accumulation of evidence towards a decision

It is well known that decisions tend to be more accurate when subjects are given longer exposure to the stimulus (see Luce (1986)). Roitman and Shadlen (2002) use a random dot stimuli experiment to study reaction times (hereafter RT). They measure activity of the neurons in the lateral intraparietal cortex (area LIP) of the monkey. Those neurons respond to visual stimuli that are the target of a voluntary eye movement and their activity is known to predict eye movement response (see Shadlen and Newsome (1996), (2001)). The authors compare two treatments, one in which subjects were exposed to the stimulus for a fixed period of time, and a second in which they were exposed to it until they reached a decision. The spike rate was found to be correlated with coherence, response and response time. More precisely, a steeper rise in the spike rate was correlated with a stronger stimulus and a shorter RT. The authors found evidence in favor of a RT model (as in Hanes and Schall (1996) or Ratcliff and Rouder (1998)) and hypothesized that LIP neurons accumulate evidence (from extrastriate visual cortex, e.g. MT neurons) over time for or against a specific response.

This evidence emphasizes that decision-making is compatible with a measure Δ that is ‘constructed’ through time from s . Information is accumulated up to a point where taking a particular action is desirable. Accumulation can be understood as a random walk toward a threshold representing the level of evidence necessary to make a decision.

Even though the analysis of RT on neurophysiological data is recent, RT distributions obtained from behavioral data have received a lot of attention in other areas. The related studies can help model neurophysiological data as well. In particular, the diffusion model (see Ratcliff (1978)) offers an appropriate framework to study simple two-choice decisions. It presupposes that decisions are made via a noisy process that accumulates information over time to one of two boundaries. The process terminates when a boundary is reached, which triggers a decision. Figure 5 illustrates the model with boundaries x and y .

[INSERT FIGURE 5]

The rate of accumulation of information is the drift rate and it reflects the quality of the information contained in the stimulus. Accumulation is noisy and in particular processes with the same drift do not always reach a boundary at the same time—which produces RT distributions—and do not always reach the same boundaries—which generates errors (see Ratcliff and McKoon (2008) for an illustration of the diffusion model using a random dot stimuli experiment on human subjects).²⁸

This model is similar in spirit to a class of models proposed in neurophysiology to fit neural activity. In a pioneering work, Hanes and Schall (1996) designed an experiment where eight targets were displayed visually to monkeys. Seven had the same color and monkeys had to detect the target with the different color by initiating a saccadic eye movement. The authors noted that the speed with which saccades were produced was variable. At the same time, neurons in the frontal eye field (an area responsible for voluntary eye movement) were found to be most active before a movement in their preferred direction. The authors showed that neurons rose to a peak firing rate very quickly and a decision process was initiated later on: neuronal firing rates continued to grow or dropped back to baseline. The authors proposed a model similar to the diffusion model to account for the behavioral and neurophysiological data.²⁹ This and further studies show that the diffusion process in the diffusion model accounts for neural activity (see Ratcliff and McKoon (2008), Gold and Shadlen (2007) and Rangel (2009) for details).

Other authors examine to what extent the behavior of populations of neurons approximates diffusion processes (see Mazurek and Shadlen (2002)). The ‘race’ model, is summarized by two neural integrators that correspond to the two alternatives and a choice is made when one of the integrators exceeds a threshold (see Luce (1986), Carpenter and Williams (1995), Shadlen and Newsome (2001)). For suitably chosen parameters, the race and diffusion models predict the same behavioral data (see Van Zandt et al. (2000) and Bogacz (2006) for details).

Overall, in this class of model, decision-making consists in simplifying D into a time-dependent measure. At each instant t , information is compiled in Δ_t . If this measure has not reached a cutoff, no decision is made. Instead, DM extracts more information from

²⁸Ratcliff and McKoon (2008) showed that the model fits the behavioral data well and in particular error responses were slower than correct responses. The main difference compared to behavioral data obtained with monkeys is that RT distributions are right skewed with humans and symmetric with monkeys.

²⁹The authors studied two specifications: a variable rate model and a variable threshold model. They also showed that regardless of the rate of increase in neuronal activity, the movement occurred when the neuron reached roughly the same firing rate, providing support for a variable rate model. See also Schall and Thompson (1999) and Schall (2003) for further studies of the various stages (perception and response preparation) that contribute to reaction times.

D and a new measure Δ_{t+1} is obtained. This process continues until a date τ when Δ_τ reaches either boundary.

Result 5 *Decision-making is consistent with a noisy process that accumulates information Δ_t over time to one of two boundaries x and y , and which terminates when a boundary is reached. Reaching x is evidence of A and triggers a while reaching y is evidence of B and triggers b .*

4.3 Summary

Taken together, these models emphasize that the information encoded by sensory neurons will be transformed into a measure to be compared to threshold point(s) triggering a decision. Yet, the models are mechanical in that they offer a description of a plausible mechanism (that fits behavioral data) but do not explain for instance why this mechanism is compatible with ‘good’ decision-making or where do the threshold(s) come from.

The purpose of the next section is to link the two sets of results reported in sections 3 and 4. One important feature of expected utility theory is *optimization*. DM chooses the option that maximizes his objective. This offers an intuitive criterion to choose between alternatives and should in principle provide foundations for the mechanical models outlined in section 4. For instance, if DM maximizes some expected payoff, thresholds in results 4 and 5 should result from optimization. Only a few articles have addressed these important issues (see Bogacz et al. (2006)). However, the studies take the mechanism as given and attempt to optimize it locally only (e.g. they optimize one bound given an objective). In particular, they do not address the reasons why optimizing is a relevant approach, why the mechanism should be a threshold mechanism and they do not discuss the objective function to maximize. In the next section, we revisit decision-making in our experiment using the economics methodology and we investigate the links between our results and the models proposed in neuroscience.

5 Economic models of decision-making

5.1 An economic diffusion model

In this section, we consider a model of information accumulation for our experiment which is reminiscent of a discrete version of the diffusion model outlined in section 4.2. The model is also reminiscent of Rustichini (2008) (and to some extent Dickhaut et al. (2009)) and captures the basic trade-offs emphasized there. Our simple model allows to encompass different experimental scenarii and derive different predictions.

The model goes as follows. The experimenter chooses state A with probability p_0 and state B with probability $1 - p_0$. These prior probabilities are common knowledge, but the realization of the choice is not known by DM. We shall assume there is a final time T at which a decision must be made, although T might be arbitrarily large. If DM decides not to make any decision in period t ($< T$), he receives a signal $\sigma_t \in \{\alpha, \beta\}$ about the true state of the world that can be used to update his beliefs. Signals α and β are imperfectly correlated with states A and B . Noise represents both the noise introduced by the experimenter and the internal noise due to cell firing variability (see Result 2). We model this in a very simple way:³⁰

$$\Pr(\alpha | A) = \Pr(\beta | B) = \theta \ (> 1/2) \quad \text{and} \quad \Pr(\alpha | B) = \Pr(\beta | A) = 1 - \theta.$$

Denote by $n_{\alpha,t}$ and $n_{\beta,t}$ the number of signals α and β received respectively up to date t . From standard probability theory, if draws of σ_t are independent, then

$$\Pr(A | n_{\alpha,t}, n_{\beta,t}) = \frac{\theta^{n_{\alpha,t} - n_{\beta,t}} p_0}{\theta^{n_{\alpha,t} - n_{\beta,t}} p_0 + (1 - \theta)^{n_{\alpha,t} - n_{\beta,t}} (1 - p_0)}.$$

Note that two different signals cancel each other out. Therefore, the relevant variable that will be used from now on is n_t ($= n_{\alpha,t} - n_{\beta,t}$), that is, the difference between the number of α signals and the number of β signals up to date t . As the agent receives one signal per period, $n_t \in \{-t, \dots, t\} \subset \mathbb{Z}$. In the notations of the previous sections, n_t represents the summary measure Δ_t of the information D encoded by the sensory system up to time t . We can define the following function:³¹

$$p(n_t) \equiv \Pr(A | n_{\alpha,t}, n_{\beta,t}) = \frac{\theta^{n_t} p_0}{\theta^{n_t} p_0 + (1 - \theta)^{n_t} (1 - p_0)}. \quad (2)$$

Recall that Result 3 indicates that Bayesian inferences can be made, that is the above calculation can be implemented in the brain. Now, according to Result 1, DM will evaluate the alternatives at their respective expected payoffs. The expected benefit of taking action a for a difference of signals n_t is

$$\pi^a(n_t) = p(n_t)G_A.$$

and the benefit of taking action b is

$$\pi^b(n_t) = (1 - p(n_t))G_B.$$

³⁰The stochastic relation between the true state and the information perceived is captured by θ and corresponds roughly to the drift in the continuous time diffusion model.

³¹Note that: (i) $p(n_t + 1) > p(n_t)$; (ii) $p(0) = p_0$; (iii) $\lim_{T \rightarrow +\infty} p(-T) = 0$; and (iv) $\lim_{T \rightarrow +\infty} p(T) = 1$.

Denote by $\gamma(n_t)$ the probability that if at some period t the difference of signals is n_t , then at $t + 1$ the difference is $n_t + 1$. Formally, $\gamma(n_t) \equiv \Pr(n_{t+1} = n_t + 1)$. We have:

$$\begin{aligned}\gamma(n_t) &= \Pr(\alpha | A)p(n_t) + \Pr(\alpha | B)(1 - p(n_t)) \\ &= (1 - \theta) + (2\theta - 1)p(n_t).\end{aligned}$$

Note that $\gamma(n_t + 1) > \gamma(n_t)$. Due to our binary signal structure, $1 - \gamma(n_t) \equiv \Pr(n_{t+1} = n_t - 1)$. Also, for any n_t , $p(n_t) = \gamma(n_t)p(n_t + 1) + (1 - \gamma(n_t))p(n_t - 1)$. Note that given the independence of signals, we can suppress time subscripts in the difference of signals.

At each date t , DM can either take action a , action b or delay the decision until date $t+1$ to receive a new signal. In case an action is taken, the corresponding reward is obtained when it applies. We assume that the representation of immediate rewards differs from the representation of future rewards, and in particular any delayed reward is discounted by the discount factor δ . This assumption is supported by a large body of research that has been reviewed in other studies (see Rangel et al. (2008) for example). An alternative model would be that attention is costly: delaying the decision to accumulate an extra piece of information induces a disutility c due to attention costs. Both models deliver the same qualitative properties and we will concentrate on the first one for the sake of simplicity.

Denote by $F_t(n)$ the value function of the agent at date $t \in \{0, \dots, T\}$ when the difference of signals is n . We have

$$F_t(n) = \max\{\pi^a(n); \pi^b(n); \delta[\gamma(n)F_{t+1}(n + 1) + (1 - \gamma(n))F_{t+1}(n - 1)]\}$$

At date t , DM chooses between the three available alternatives corresponding to the three respective payoffs in the above value functions, and selects the alternative yielding the highest payoff. His problem is to determine under which condition it is optimal to choose any of those alternatives. In other words, we want to determine the range of the values n for which it is best to choose a , b or to delay the decision one period. Solving this problem is equivalent to maximizing expected utility in our setting. Importantly, our methodology does not assume decision-making is made via thresholds, rather the model delivers threshold values. Furthermore, given our optimization process, these threshold are optimal, that is they are set in such a way DM achieves highest payoffs. Let us note that Gold and Shadlen (2002) proposed a diffusion model in which subjects choose x so as to maximize the reward obtained in the experiment. This approach has been followed by others (see Bogacz (2007) for a review). However, all these studies presuppose that the mechanism is a threshold-mechanism or impose arbitrary bounds.

Proposition 1 *When a decision is due before time T , DM continues to accumulate information at date $t \leq T$ if and only if $n_t \in (n_t^{b*}, n_t^{a*})$. DM stops if $n_t = n_t^{a*}$ in which case he undertakes a , or if $n_t = n_t^{b*}$ in which case he undertakes b . The thresholds have the following properties:*

- (i) *the lower threshold increases over time, i.e. $n_t^{b*} < n_{t+1}^{b*}$ for all t*
- (ii) *the higher threshold decreases over time, i.e. $n_t^{a*} > n_{t+1}^{a*}$ for all t*
- (iii) *at the last period $n_T^{b*} = n_T^{a*} = \tilde{n}$ where \tilde{n} is such that $\pi^a(n) = \pi^b(n)$.*

For all t , n_t^{a} is decreasing in G_A and increasing in G_B , n_t^{b*} is increasing in G_B and decreasing in G_A . Furthermore, for all $t < T$, n_t^{a*} is increasing in δ and n_t^{b*} is decreasing in δ .*

Proof: See Appendix.

At each date, there are two optimal thresholds. Action a is triggered if n reaches n_t^{a*} , while action b is triggered if n reaches n_t^{b*} . If n lies in (n_t^{b*}, n_t^{a*}) , DM waits one more period. In the last period, DM must make a decision and the two thresholds coincide. It corresponds to the difference of signals such that the payoffs of taking either action are the same. In period $T - 1$, DM should wait if his difference of signal is \tilde{n} . Delaying one period will provide him with better evidence about the state. In other words, there is a positive option value of waiting. The threshold above which it is optimal to choose a is strictly above \tilde{n} whereas the threshold below which it is optimal to choose b is strictly below \tilde{n} . The option value of waiting is as large as the horizon is large and as time passes, it becomes less beneficial to wait.³²

Naturally, the magnitude of (n_t^{b*}, n_t^{a*}) – which we will call the ‘waiting region’ – affects the probability of reaching a decision sooner or later. That is, the model makes also predictions in terms of reaction times. The bigger the reward of an alternative, the higher the incentives to bias the decision towards that alternative. Anticipating there is a high chance to choose it in the future, DM might as well choose it earlier. Formally, n_t^{a*} decreases when G_A increases, and n_t^{b*} increases when G_B increases, making DM more likely to choose early when rewards are high. Also, the higher the discount factor, the higher the option value of waiting. Therefore, it is optimal to set n_t^{b*} as low as δ is high and n_t^{a*} as high as δ is high. The waiting region being enlarged, the decision is delayed more often.

[INSERT FIGURE 6]

³²These results are reminiscent of those obtained in the literature on investment and uncertainty (see Dixit and Pindyck (1994))

Our analysis predicts that the mechanical diffusion model has an optimization model counterpart. This is in itself already important. An added benefit of the optimization approach is the possibility to perform comparative statics about the object of interest. We show first that the optimal thresholds vary over time. Second, they vary in a precise way. As time passes, it becomes more urgent to make a decision. In particular, DM might wait in the early stages for a given value n but act upon it in the later stages. Our results have also a few instructive corollaries.

Corollary 1 *If the individual has no time constraint ($T = +\infty$), then $n_t^{a*} = n^{a*}$ and $n_t^{b*} = n^{b*}$ for all t . Furthermore,*

- (i) $n^{a*} = n^{b*} = \tilde{n}$ when $\delta \rightarrow 0$,
- (ii) $n^{a*} \rightarrow +\infty$ and $n^{b*} \rightarrow -\infty$ when $\delta \rightarrow 1$.

Last, the error rate in taking either action increases with the reward of that action, and the reaction time in taking either action decreases with the reward of that action.

Proof: See Appendix.

In the absence of time constraint, the option value of waiting is the same at all dates. Therefore, the optimal thresholds need not be adjusted over time. In that case, *and in that case only*, boundaries are time invariant at the optimum. Interestingly, boundaries are also affected by the size of the discount factor. If DM does not value the future (e.g. cannot represent future rewards), each date is terminal from his perspective. The only relevant threshold is \tilde{n} . At equilibrium, he will make a decision right away and never delay it. On the other extreme, if DM does not distinguish between present and future (i.e., DM is infinitely patient), the option value of waiting is positive for any value n and at any point in time. DM will delay the decision indefinitely. Taken together, these results suggest that the diffusion model with fixed boundaries can be rationalized in experiments with no (implicit or explicit) time constraint and for intermediate values of time discounting. Obviously, data collected from different subjects should fit different parameter values. Impatient (respectively patient) subjects are best represented by a model with a narrow (respectively wide) waiting region (n^{b*}, n^{a*}). These results are summarized in figure 7.

[INSERT FIGURE 7]

Last, it is interesting to note that the model makes clear predictions in terms of error rates. The probability of choosing action a while the true state is A is equal to the probability of hitting n^{a*} before n^{b*} . Given those thresholds are modulated by the magnitude of rewards, the error rates depend on G_A and G_B as well. We know already that it is efficient to lower the threshold when the reward of action a increases. Therefore, this action will be taken more often, also wrongly. Also, when G_A increases (in which case both n_a^* and n_b^*

decrease) the reaction time of taking action a decreases while the reaction time in taking action b increases.

The results in this section validates the overall mechanical layout of the diffusion model. However, they suggest that a model with fixed and arbitrary boundaries might not be the most accurate representation of decision-making. Rather, boundaries should be time-dependent and derived from a utility based optimization problem. The above results also indicate that discounting is a necessary condition to have two interior threshold values. If DM is patient (or simply patient over the length of the experiment), it is optimal to accumulate as much information as possible and make a decision at the last minute.³³ In a recent study, Resulaj et al. (2009) showed that decision making is consistent with a phase of accumulation of information towards a boundary triggering an initial decision, followed by a phase in which the brain exploits information to either reaffirm or reverse the initial decision. Even though the setting of the experiment is different from the model proposed here, the evidence is consistent with one prediction of our model: thresholds may change over time.

5.2 A single threshold model of information processing

We now investigate a variant in which DM obtains the reward at date T independently of the time at which he makes a decision. In such experiment, discounting does not play a role. It is optimal to wait until the last period to make a decision, that is $n_t^{a*} = +\infty$ and $n_t^{b*} = -\infty$ for all $t < T$ and $n_T^{a*} = n_T^{b*} = \tilde{n}$. Said differently, the waiting region is unbounded up to date T and collapses into a single decision threshold at T . This model is closer in spirit to the models reviewed in section 4.1. Our optimal threshold model suggests that Δ^* in Result 4 should be consistent with \tilde{n} .

Proposition 2 *When rewards are delayed to the last period, DM should accumulate information until the last minute and select action a if $n > \tilde{n}$ and action b if $n < \tilde{n}$. The optimal threshold \tilde{n} satisfies*

$$\frac{\theta^{\tilde{n}}}{(1-\theta)^{\tilde{n}}} = \frac{G_B}{G_A} \frac{1-p_0}{p_0}$$

*It is decreasing in p_0 , increasing in G_A and decreasing in G_B .*³⁴

Proof: See Appendix.

³³This holds because waiting is costly. In our setting, the cost of waiting is indirect and due to the discount factor. Alternatively, and as noted earlier, we could assume DM faces a direct cost for obtaining extra evidence (e.g. an attention cost). The results obtained here would hold qualitatively.

³⁴Equivalently, the optimal threshold is $\tilde{n} = \log(\frac{G_B}{G_A} \frac{1-p_0}{p_0}) / \log(\frac{\theta}{1-\theta})$.

The model is formally equivalent to a simple version of Brocas and Carrillo (2011) and the result has two implications. First, the optimal single threshold is set in such a way that existing beliefs are likely to be confirmed. Suppose that the prior belief p_0 is such that it is a priori optimal to take action a (i.e. $p_0 G_A > (1 - p_0) G_B$). New evidence will either confirm a is the correct decision or provide evidence b should be taken instead. If the threshold is high, DM must receive strong support of a to stick to that decision. If the threshold is surpassed, DM takes action a with a high confidence level. If it is not surpassed however, DM still believes a was a good option but he implements b . Overall, DM is confident in his decision only in one scenario. Instead, if the threshold is low, DM must collect a lot of evidence in favor of b to switch to that decision. If it is surpassed, DM sticks to his initial belief as in the previous case. If the threshold is not reached however, DM is now convinced that the state is B and takes action b confidently. DM ends up taking different actions depending on the information received but, in both cases, he is confident about his choice. Overall, the optimal threshold balances the belief in favor of A conditional on the threshold being surpassed and the belief in favor of B conditional on the threshold not being reached.

Second, the optimal single threshold is affected by the relative rewards. If the loss of taking the wrong action when the true state is A increases, the threshold decreases, so it will be surpassed with higher probability. In equilibrium, the most costly mistakes are most likely to be avoided.

Note that the optimal threshold is in line with the proposal made in Gold and Shadlen (2001). Recall that they proposed to compare $\log \frac{P(D|A)}{P(D|B)}$ to a cutoff $\Delta^*(p_0, G_A, G_B)$. In our model, the information accumulated D is summarized into n , and our optimal threshold is such that

$$\frac{\theta^{\tilde{n}}}{(1 - \theta)^{\tilde{n}}} \equiv \frac{P(D|A)}{P(D|B)} = \frac{G_B}{G_A} \frac{1 - p_0}{p_0}.$$

Therefore $\Delta^* = \log \left(\frac{G_B}{G_A} \frac{1 - p_0}{p_0} \right)$.

Interestingly, the optimal threshold affects the collection of information in such a way that DM is more likely to end up taking action a when p_0 is originally high. This implies that DM will tend to confirm his prior. Also, DM is more likely to end up taking action a if G_A is relatively more attractive. Therefore, DM acts as if his ex-post beliefs are shaped by the outcomes. These two conclusions are reminiscent of Brocas and Carrillo (2011).³⁵ In this analysis, the authors show that these two fundamental biases are likely the source of some well documented behavioral anomalies. In particular, the tendency to

³⁵The authors consider a more general model in a non parametric framework (their model requires only the monotone likelihood ratio property to hold).

reinforce priors could be the reason why DM polarize when exposed to the same mixed evidence, or why the order in which evidence is produced affect decisions. Also, the fact that outcomes affect posterior beliefs explains that two DM with the same prior and facing the same information may end up disagreeing. More generally, given the impact of outcomes on posterior beliefs, decisions are best represented by rank-dependent or security-potential/aspiration type of utility, and the analysis therefore provides support for some form of relaxation of expected utility theory. It also explains why the same DM may take decisions in apparent contradictions with each other. For instance, if an individual is most afraid of the damages of an earthquake, he will interpret information in such a way he may end up assigning a large probability to that event and therefore purchase a very expensive insurance. If at the same time, he does not care as much about his health, he will tend to believe he is not likely to have health issues and contract a minimal health insurance. These decisions are internally consistent but might appear in contradiction to an outside observer. If behavior is required to fit a standard expected utility model, the model will also pick different levels of risk aversion in the different situations.

6 Conclusion

While the overlap between the paradigms in economics and neuroscience have already been noted,³⁶ an open debate is still to determine how exactly can each field contribute to the understanding of the other. This article aims at providing a few reasons why economists should pay attention to findings obtained in the Brain Sciences. We know from a large body of research that humans (e.g. from the social psychology and experimental economics literatures) and animals (e.g. from the behavioral ecology literature) often make decisions in ways that are efficient.³⁷ However, and rather systematically, DM behave in a way that is not consistent with the previous view. The causes of departures are still unknown. We believe that a better understanding of brain functions and processes should help shedding light on systematic anomalies. We can derive two strong basic messages from the evidence reviewed here and the models we built.

First, economic decision making paradigms are roughly consistent with brain functioning. Neuronal processes interpret information and make decisions in a way that is consistent with the standard paradigm in economics: both ‘expected utility-like’ theory (Result 1) and ‘Bayesian-like’ information acquisition theory (Result 3) are relevant hy-

³⁶See for instance Glimcher et al. (2005).

³⁷Behavioral ecologists favor the hypothesis that evolution pushes animals toward efficient decision-making. They build models based on economic theory to account for the computations that animals perform. See Glimcher (2002).

potheses.

Second, the standard paradigm in Economics takes for granted a large number of features that may distort our understanding of decision-making. Indeed, some important ‘details’ may cause differences between the predictions of these theories and observed choices. Importantly, the brain accumulates information until there is enough information supporting a behavioral response. This means it uses an economical process in which the information just necessary to make a decision is taken into account, and decision-making is best represented by threshold mechanisms (Results 4 and 5). Recommendations are collected from various entities (neurons in the sensory systems) and aggregated into a single measure that is then compared to threshold(s) to formulate a decision. We know from the social choice literature that any aggregation rule has a preferred direction and ultimately favors some features. Therefore, the aggregation of information throughout layers of populations of neurons will yield distortions of the signal, as some features will be emphasized and others de-emphasized.

The models we proposed in section 5 combine these pieces of evidence to make predictions about behavior. Recognizing the validity of expected utility theory leads to one fundamental conclusion: DM acts as if he is optimizing his objective. This provides foundations for using an optimization approach. This is a drastic departure from the existing mechanical models used in Neuroscience. At the same time, the optimization approach can provide foundations for such models. The analysis provides also a few reasons why neuroscientists should pay attention to economic modeling. In particular, the approach is rich as it derives a solution (an optimal threshold mechanism) as a function of primitives (variability, physiological constraints, constraints set by the experimenter, etc.), rather than assuming the solution. Propositions 1 and 2 are examples of the methodology. Not only they establish that threshold models are optimal mechanisms in the environments considered, but they also suggest how thresholds should be set, and eventually give clear predictions about behavior. Such predictions are in principle testable with neural and behavioral data. We believe this constitutes an interesting alley of research in Neuroscience.

Other features reviewed here may be the source of other anomalies. For instance, given neuronal variability, the notions of ‘complete information’, ‘perfect signal’ or even ‘objective probabilities’ might not be pertinent in the brain. Basically, the brain is typically ignorant or uninformed before a signal is sent and it decodes the signal through variable processes. This may account for framing effects in situations where information is presumably disclosed or verifiable. For instance, revisiting experimental results in complete information settings with this in mind may help understanding inefficient behavior.

Last, from a normative perspective, it is useful to decompose the ingredients of a de-

cision problem into utilities and probabilities. However, it is important to keep in mind that both concepts are represented in a similar way in the brain. Being offered a lottery, announced a weather forecast or showed a cookie triggers an informational process: neurons detect the characteristics of the item offered, fire in a probabilistic manner, pass the information along, part of the information is de-emphasized, this information is aggregated and processed to obtain a response, and last the response is implemented. Actual decision making results from this underlying informational process. Understanding it better should shed light in a meaningful way on decision theoretic paradigms. Concretely, we may want to think about new theories in which brain processes are the primitives (or axioms) to derive utility representations. The constant effort to collect evidence should allow to build such new theories in the future.

Appendix

Proof of Proposition 1

Let \tilde{n} be such that $\pi^a(n) = \pi^b(n)$, then the previous value function can be rewritten as:

$$F_t(n) = \begin{cases} \max \left\{ \pi^a(n); \delta[\gamma(n)F_{t+1}(n+1) + (1-\gamma(n))F_{t+1}(n-1)] \right\} & n > \tilde{n} \\ \max \left\{ \pi^b(n); \delta[\gamma(n)F_{t+1}(n+1) + (1-\gamma(n))F_{t+1}(n-1)] \right\} & n < \tilde{n} \end{cases}$$

When $n > \tilde{n}$, the optimization problem is equivalent to

$$H_t^a(n) = \max\{0; -(1-\delta)\pi^a(n) + \delta[\gamma(n)H_{t+1}^a(n+1) + (1-\gamma(n))H_{t+1}^a(n-1)]\}$$

where $H_t^a(n) = F_t(n) - \pi^a(n)$. Similarly, when $n < \tilde{n}$, the optimization problem is equivalent to

$$H_t^b(n) = \max\{0; -(1-\delta)\pi^b(n) + \delta[\gamma(n)H_{t+1}^b(n+1) + (1-\gamma(n))H_{t+1}^b(n-1)]\}$$

where $H_t^b(n) = F_t(n) - \pi^b(n)$.

At time T , it is optimal to choose action a over b if $n > n_T^* = \tilde{n}$.

At time $T-1$, there are two cases. If $n > \tilde{n}$, the optimization problem is equivalent to

$$H_{T-1}^a(n) = \max\{0; -(1-\delta)\pi^a(n) + \delta[\gamma(n)H_T^a(n+1) + (1-\gamma(n))H_T^a(n-1)]\}$$

where $H_T^a(n) = F_T(n) - \pi^a(n)$. Note that $\pi^a(n)$ is strictly increasing in n , and therefore $H_T^a(n)$ is non-increasing in n . Given $\gamma'(n) < 0$, then we also have

$$\frac{\partial}{\partial n} \left[\gamma(n)H_T^a(n+1) + (1-\gamma(n))H_T^a(n-1) \right] < 0 \quad (3)$$

and overall the (r.h.s.) of $H_{T-1}^a(n)$ is strictly decreasing in n and the value function $H_{T-1}^a(n)$ is non increasing in n . As a consequence, there exists n_{T-1}^{a*} that satisfies

$$\pi^a(n_{T-1}^{a*}) = \delta \left[\gamma(n_{T-1}^{a*})F_T(n_{T-1}^{a*}+1) + (1-\gamma(n_{T-1}^{a*}))F_T(n_{T-1}^{a*}-1) \right] > 0.$$

Therefore, $n_{T-1}^{a*} > n_T^*$.

If $n < \tilde{n}$, the optimization problem is equivalent to

$$H_{T-1}^b(n) = \max\{0; -(1-\delta)\pi^b(n) + \delta[\gamma(n)H_T^b(n+1) + (1-\gamma(n))H_T^b(n-1)]\}$$

where $H_T^b(n) = F_T(n) - \pi^b(n)$. Note that $\pi^b(n)$ is strictly decreasing in n , and therefore $H_T^b(n)$ is non-decreasing in n . Given $\gamma'(n) < 0$, then we also have

$$\frac{\partial}{\partial n} \left[\gamma(n)H_T^b(n+1) + (1-\gamma(n))H_T^b(n-1) \right] > 0 \quad (4)$$

and overall the (r.h.s.) of $H_{T-1}^b(n)$ is strictly increasing in n and the value function $H_{T-1}^b(n)$ is non decreasing in n . As a consequence, there exists n_{T-1}^{b*} that satisfies

$$\pi^b(n_{T-1}^{b*}) = \delta \left[\gamma(n_{T-1}^{b*})F_T(n_{T-1}^{b*} + 1) + (1-\gamma(n_{T-1}^{b*}))F_T(n_{T-1}^{b*} - 1) \right] > 0.$$

Therefore, $n_{T-1}^{b*} < n_T^*$.

To characterize the optimal cutoffs at times $t < T$, we use a recursive method. Suppose that the following properties hold at date t when $n > \tilde{n}$:

- (A1) $H_t^a(n)$ is non increasing in n ;
- (A2) $H_t^a(n) \geq H_{t+1}^a(n)$.

$$\begin{aligned} H_t^a(n) &= \max \left\{ 0; -(1-\delta)\pi^a(n) + \delta[\gamma(n)H_{t+1}^a(n+1) + (1-\gamma(n))H_{t+1}^a(n-1)] \right\}; \\ H_{t-1}^a(n) &= \max \left\{ 0; -(1-\delta)\pi^a(n) + \delta[\gamma(n)H_t^a(n+1) + (1-\gamma(n))H_t^a(n-1)] \right\}. \end{aligned}$$

Combining (A1) and the fact that $\gamma'(n) > 0$, we get that the r.h.s. of $H_{t-1}^a(n)$ is decreasing in n . Therefore, if $H_t^a(n)$ is non-increasing in n , $H_{t-1}^a(n)$ is also non-increasing in n . Moreover, there exists a unique cutoff n_{t-1}^{a*} above which the agent takes action a at date $t-1$. In addition, given (A2), it is clear that

$$\gamma(n) \left[H_t^a(n+1) - H_{t+1}^a(n+1) \right] + (1-\gamma(n)) \left[H_t^a(n-1) - H_{t+1}^a(n-1) \right] \geq 0.$$

and therefore, the r.h.s. of $H_{t-1}^a(n)$ is greater than the r.h.s. of $H_t^a(n)$, which is sufficient to prove that $H_{t-1}^a(n) \geq H_t^a(n)$. Overall, we have shown that if (A1) and (A2) hold at date t , then they also hold at date $t-1$, which completes the recursion. This in turn implies that $n_{t-1}^{a*} > n_t^{a*}$ for all t .

Suppose that the following properties hold at date t when $n < \tilde{n}$:

- (A1') $H_t^b(n)$ is non decreasing in n ;
- (A2') $H_t^b(n) \leq H_{t+1}^b(n)$.

$$\begin{aligned} H_t^b(n) &= \max \left\{ 0; -(1-\delta)\pi^b(n) + \delta[\gamma(n)H_{t+1}^b(n+1) + (1-\gamma(n))H_{t+1}^b(n-1)] \right\}; \\ H_{t-1}^b(n) &= \max \left\{ 0; -(1-\delta)\pi^b(n) + \delta[\gamma(n)H_t^b(n+1) + (1-\gamma(n))H_t^b(n-1)] \right\}. \end{aligned}$$

Combining **(A1')** and the fact that $\gamma'(n) > 0$, we get that the r.h.s. of $H_{t-1}^b(n)$ is increasing in n . Therefore, if $H_t^b(n)$ is non-decreasing in n , $H_{t-1}^b(n)$ is also non-decreasing in n . Moreover, there exists a unique cutoff n_{t-1}^{b*} below which the agent takes action b at date $t-1$. In addition, given **(A2')**, it is clear that

$$\gamma(n) \left[H_t^b(n+1) - H_{t+1}^b(n+1) \right] + (1 - \gamma(n)) \left[H_t^b(n-1) - H_{t+1}^b(n-1) \right] \geq 0.$$

and therefore, the r.h.s. of $H_{t-1}^b(n)$ is smaller than the r.h.s. of $H_t^b(n)$, which is sufficient to prove that $H_{t-1}^b(n) \leq H_t^b(n)$. Overall, we have shown that if **(A1')** and **(A2')** hold at date t , then they also hold at date $t-1$, which completes the recursion. This in turn implies that $n_{t-1}^{b*} < n_t^{b*}$ for all t .

We can also use a recursion to show that n_t^{a*} is decreasing in G_A . It is easy to see that \tilde{n} is decreasing in G_A . At date $T-1$, $\pi^a(n)$ is increasing in G_A and the r.h.s. of $H_{T-1}^a(n)$ is decreasing in G_A . Therefore $H_{T-1}^a(n)$ is non increasing in G_A and n_{T-1}^{a*} is decreasing in G_A . Assuming that $H_t^a(n)$ is non increasing in G_A , it is easy to show that $H_{t-1}^a(n)$ is also non increasing in G_A . Therefore, n_t^{a*} is decreasing in G_A for all t . Also, \tilde{n} is increasing in G_B . At date $T-1$, the r.h.s. of $H_{T-1}^a(n)$ is increasing in G_B . Therefore $H_{T-1}^a(n)$ is non decreasing in G_B and n_{T-1}^{a*} is increasing in G_B . Assuming that $H_t^a(n)$ is non decreasing in G_B , it is easy to show that $H_{t-1}^a(n)$ is also non decreasing in G_B . Therefore, n_t^{a*} is increasing in G_B for all t .

A similar recursion applies to show that n_t^{b*} is increasing in G_B and decreasing in G_A .

Last, at date $T-1$, the r.h.s. of $H_{T-1}^a(n)$ is increasing in δ . Therefore $H_{T-1}^a(n)$ is non decreasing in δ and n_{T-1}^{a*} is increasing in δ . Assuming that $H_t^a(n)$ is non decreasing in δ , it is easy to show that $H_{t-1}^a(n)$ is also non decreasing in δ . Therefore, n_t^{a*} is increasing in δ for all t . A similar recursion applies to show that n_t^{b*} is decreasing in δ .

Proof of Corollary 1

When $T \rightarrow \infty$, the problem is recursive and the value function is:

$$F(n) = \begin{cases} \max \left\{ \pi^a(n); \delta[\gamma(n)F(n+1) + (1-\gamma(n))F(n-1)] \right\} & n > \tilde{n} \\ \max \left\{ \pi^b(n); \delta[\gamma(n)F(n+1) + (1-\gamma(n))F(n-1)] \right\} & n < \tilde{n} \end{cases}$$

Therefore, the thresholds are the same for all t . When $\delta = 0$,

$$F(n) = \max\{\pi_a(n), \pi_b(n); 0\}$$

DM chooses in the first period based on one single threshold: \tilde{n} . When $\delta = 1$, the r.h.s. of the two equations are positive and it is best to wait for all n .

The probability of taking action a when the true state is B is equal to the probability $q(G_A, G_B)$ of reaching n^{a*} before n^{b*} . It has been shown in Brocas and Carrillo (2007) that this probability is simply:

$$q(G_A, G_B) = \frac{p_0 - p(n^{a*})}{p(n^{a*}) - p(n^{b*})} \frac{1 - p(n^{a*})}{1 - p_0}$$

Therefore

$$\frac{\partial q}{\partial G_A} \equiv -\frac{\partial n^{b*}}{\partial G_A} p'(n^{b*}) (p(n^{a*}) - p_0) (1 - p(n^{a*})) - \frac{\partial n^{a*}}{\partial G_A} p'(n^{b*}) (p_0 - p(n^{b*})) (1 - p(n^{b*})) > 0$$

$$\frac{\partial q}{\partial G_B} \equiv -\frac{\partial n^{b*}}{\partial G_B} p'(n^{b*}) (p(n^{a*}) - p_0) (1 - p(n^{a*})) - \frac{\partial n^{a*}}{\partial G_B} p'(n^{b*}) (p_0 - p(n^{b*})) (1 - p(n^{b*})) < 0$$

Last denote $n^{a*}(G_A)$ the threshold when the benefit of taking action a is G_A . Consider G_A and $G'_A > G_A$ such that $n^{a*}(G'_A) = n^{a*}(G_A) - 1$. Consider a path that leads to take action a at some date $\tau < t$ under G'_A (i.e. that reach $n^{a*}(G_A) - 1$) may now not lead to take action a before date t . Therefore the reaction time in taking action a increases when G_A decreases. A similar argument holds for G_B and action b .

Proof of Proposition 2

The comparative statics on \tilde{n} are obtained by differentiating equation (2).

Probability




Figure 1: D represents the spike rate when the stimulus s signals the true state S (true direction) and the curve plots the likelihood of obtaining response D .

spike counts




Figure 2a: Each dot represents the number of spikes obtained for a neuron with preferred direction A in a trial where the true direction is S . Values are recorded for all trials.

spike counts




Figure 2b: Similar illustration with more than two states. The preferred direction of the neuron is S_p

spike counts




Figure 3a: Population vector estimator

spike counts




Figure 3b: Maximum likelihood estimator (\hat{S} and \tilde{S} may not coincide)




Figure 4: Psychometric function.




Figure 5: Evidence is accumulated through time (r is the drift) until hitting the boundary x triggering a or the boundary y triggering b .




Figure 6: Optimal diffusion model with horizon T .




Figure 7: Optimal diffusion model with infinite horizon.

References

1. Abbott, L.F., 1994. Decoding neuronal firing and modeling neural networks. *Quarterly Review of Biophysics* 27, 291-331.
2. Albright, T., 1984. Direction and Orientation Selectivity of Neurons in Visual Area MT of the Macaque. *Journal of Neurophysiology* 52, 1106-1130.
3. Atick, J.J., and A.N. Redlich. 1992. What Does the Retina Know about Natural Scenes?. *Neural Computations* 4, 196-210.
4. Barlow, H., 1961. Possible principles underlying the transformation of sensory messages. in *Sensory Communication*. MIT Press.
5. Barlow, H., 2001. Redundancy Reduction Revisited. *Network: Computation in Neural Systems* 12, 241-253.
6. Bogacz, R. E., Brown, J., Moehlis, and J.D. Cohen, 2006. The physics of optimal decision-making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review* 113, 700-765.
7. Bogacz R., 2007. Optimal decision-making theories: linking neurobiology with behaviour. *Trends in Cognitive Sciences* 11, 118-125.
8. Britten K.H., M.N. Shadlen, W.T. Newsome and J.A. Movshon, 1992. The analysis of visual Motion: A Comparison of Neuronal and Psychophysical performance. *Journal of Neuroscience* 12, 4745-4765.
9. Brocas I. and J.D. Carrillo, 2007. Influence through ignorance. *RAND Journal of Economics*, 38 (4), 931-947.
10. Brocas I. and J.D. Carrillo, 2011. From perception to action: an economic model of brain processes. *Games and Economic Behavior*, forthcoming.
11. Carpenter R.H.S. and M.L.L. Williams, 1995. Neural computation of log likelihood in the control of saccadic eye movements. *Nature* 377, 59-62.
12. Cowan W.M., T.C. Sudhof and C.F. Stevens, 2003. *Synapses*. The Johns Hopkins University Press.
13. Deneve, S., Latham, P., and A. Pouget, 1999. Reading Population Codes: a Neural Implementation of Ideal Observers. *Nature Neuroscience* 2(8), 740-745.
14. Dickhaut, J., Rustichini, A., and V. Smith, 2009. A neuroeconomic theory of the decision process. *Proceedings of the National Academy of Sciences*. 106(52), 22145-22150.

15. Ditterich, J., Mazurek, M., and M. Shadlen, 2003. Microstimulation of Visual Cortex Affects the Speed of Perceptual Decisions. *Nature Neuroscience* 6(8), 891-898.
 16. Dixit A.K. and R.S. Pindyck, 1994. *Investment under uncertainty*. Princeton University Press.
 17. Field D.J., 1994. What is the goal of sensory coding?. *Neural Computation* 6, 559-601.
 18. Georgopoulos, A.P., Kettner, R.E., and A. Schwartz, 1986. Neural population coding of movement directions. *Science*, 243, 234-236.
 19. Gerstner W. and W.M. Kistler 2002. *Spiking Neuron Models. Single Neurons, Populations, Plasticity*. Cambridge University Press.
 20. Gielen, C.C.A.M., Hesselms, G.H.F.M., and P.I.M. Johannesma, 1988. Sensory interpretation of neural activity patterns. *Mathematical Biosciences*, 88, 13-35.
 21. Glimcher, P., Dorris, M. and H Bayer 2005. Physiological Utility Theory and the Neuroeconomics of Choices. *Games and Economic Behavior*, 52, 213-256.
 22. Glimcher, P., 2002. Decisions, decisions, decisions: choosing a biological science of choice. *Neuron* 36, 323-332.
 23. Glimcher, P and A. Rustichini, 2004. Neuroeconomics: the Consilience of Brain and Decision. *Science* 306, 447-452.
 24. Gold, J.I and M.N. Shadlen, 2001. Neural Computations that Underlie Decisions about Sensory Stimuli. *Trends in Cognitive Sciences* 404, 390-394.
 25. Gold, J.I and M.N. Shadlen, 2002. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299-308.
 26. Gold, J.I and M.N. Shadlen, 2007. The neural basis of decision-making. *The Annual review of Neuroscience* 30, 535-574.
- XXXXXXXXXX
27. Hanes, D.P. and J.D. Schall, 1996. Neural Control of Voluntary Movement Initiation. *Science* 247, 427-430.
 28. Hanks T.D., M.E. Mazurek, R. Kiani, E. Hopp and M.N. Shadlen, 2011. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *The Journal of Neuroscience* 31, 6339-6352.
 29. Hebb D.O., 1949. *The Organization of Behavior: a Neurophysiological Theory*. J. Wiley, New-York.

30. Hodgkin A.L., Huxley AF, Katz B. 1952. A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology* 117, 500-544.
31. Kable J.W. and P.W. Glimcher, 2009. The neurobiology of decision: consensus and controversy. *Neuron* 63, 733-745.
32. Klyachko V. and C. Stevens, 2006. Excitatory and Feed-Forward Inhibitory Hippocampal Synapses Work Synergistically as an Adaptive Filter of Natural Spike Trains. *PLoS Biology* 4(7), 1187-1200.
33. Lau B., and P.W. Glimcher, 2008. Value representations in the primate striatum during matching behavior. *Neuron* 58, 451-463.
34. Laughlin S.B., R.R. van Steveninck and J.C. Anderson, 1998. The metabolic cost of neural information. *Nature Neuroscience* , 1, 36 - 41.
35. Lewen G.D., Bialek W and R.R. de Ruyter van Steveninck, 2001. Neural Coding of Naturalistic Motion Stimuli. *Network* 12, 317-329.
36. Luce R.D., 1986. Response times: their role in inferring elementary mental organization. Oxford University Press.
37. Ma, W.J., Beck, J.M., Latham, P.E. and A. Pouget, 2006. Bayesian Inference with Probabilistic Population Codes. *Nature Neuroscience* 9(11), 1432-1438.
38. Mazurek M.E. and M.N. Shadlen, 2002. Limits to the temporal fidelity of cortical spike rate signals. *Nature Neuroscience* 5, 463-471.
39. McCarthy S.T. and W.J. Owen, 1996. Preferential Representation of Natural Scenes in the Salamander retina. *Investigative Ophthalmology and Visual Science* 37, S674.
40. McCoy A. and M.L. Platt, 2005. Expectations and outcomes: decision-making in the primate brain. *Journal of Comparative Physiology A* , 191, 201-211.
41. Miller E.K. and J.D. Cohen, 2001. An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience* 24, 167-202.
42. Newsome W.T., Britten K.H. and J.A. Movshon, 1989. Neuronal correlates of a perceptual decision. *Nature* 341, 52-54.
43. Nicholls J.G., A.R. Martin, B.G. Wallace, P.A. Fuchs, 2001. *From Neuron to Brain: A Cellular and Molecular Approach to the Function of the Nervous System*. Sinauer Associates.
44. Nichols, J. and W. Newsome, 2002. Middle Temporal Visual Microstimulation Influences Veridical Judgments of Motion Direction. *The Journal of Neuroscience* 22(21), 9530-9540.

45. Padoa-Schioppa, C. and J.A. Assad 2006. Neurons in Orbitofrontal Cortex Encode Economic Value. *Nature* 441, 223-226.
46. Padoa-Schioppa, C. and J.A. Assad, 2008. The Representation of Economic Value in the Orbitofrontal Cortex is Invariant for Changes of Menu. *Nature Neuroscience* 11(1), 95-102.
47. Phelps E.A., 2009. The study of emotions in neuroeconomics. In P.W. Glimcher, C.F. Camerer, E. Fehr and R.A. Poldrack (Eds.). *Neuroeconomics: Decision Making and the Brain*. Elsevier.
48. Platt, M.L. and P.W. Glimche, 1999. Neural Correlates of Decision Variables in Parietal Cortex , *Nature* 400, 233-238.
49. Rangel A., C. Camerer and R. Montague, 2008. A framework for studying the neurobiology of value-based decision-making. *Nature Reviews - Neuroscience* 9, 545-556.
50. Rangel A., 2009. The computation and comparison of value in goal-directed choice. In P.W. Glimcher, C.F. Camerer, E. Fehr and R.A. Poldrack (Eds.). *Neuroeconomics: Decision making and the Brain*. Elsevier.
51. Ratcliff R., 1978. A theory of memory retrieval. *Psychological Review* 85(2), 59-108.
52. Ratcliff R. and Rouder J.N., 1998. Modeling response times for two-choice decisions. *Psychological Science* 9, 347-356.
53. Ratcliff R. and G. McKoon, 2008. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation* 20, 873-922.
54. Resulaj A., R. Kiani, D.M. Wolpert and M.N. Shadlen, 2009. Changes of mind in decision-making. *Nature* 461, 263-268.
55. Roitman J.D. and M.N. Shadlen, 2002. Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task. *Nature Neuroscience* 22, 9475-9489.
56. Rustichini A., 2008. Neuroeconomics: formal models of decision-making and cognitive neuroscience. In P.W. Glimcher, C.F. Camerer, E. Fehr and R.A. Poldrack. *Neuroeconomics: Decision Making and the Brain*. Elsevier.
57. Samejima K., 2005. Representation of action-specific reward values in the striatum. *Science* 310, 1337-1340.
58. Schall J.D. and K.G. Thompson, 1999. Neural selection and control of visually guided eye movements. *Annual Review of Neuroscience* 22, 241-259.

59. Schneidman, E., W. Bialek, and M. J. Berry, 2003. Synergy, Redundancy, and Independence in Population Codes. *The Journal of Neuroscience*. 23(37), 11539-11553.
60. Schall J.D. 2003. Neural correlates of decision processes: neural and mental chronometry. *Current Opinion in Neurobiology* 13, 182-186.
61. Schultz W., 2006. Behavioral theories and the neurophysiology of reward. *Annual review of Psychology* 57, 87-115.
62. Shadlen, M.N., Britten, K.H., Newsome, W.T., and J.A. Movshon, 1996. A Computational Analysis of the Relationship between Neuronal and Behavioral Responses to Visual Motion. *Journal of Neuroscience* 16, 1486-1510.
63. Shadlen M.N. and W.T. Newsome, 1996. Motion Perception: seeing and deciding. *Proceedings of the National Academy of Science* 93, 628-633.
64. Shadlen M.N. and W.T. Newsome, 2001. Neural Basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology* 86, 1916-1936.
65. Shannon, C.E., 1948. A Mathematical Theory of Communication. *Bell System Technical Journal* 27, 379-423 and 623-656.
66. Shepherd G.M., 2003. *The Synaptic Organization of the Brain*. Oxford University Press.
67. Simoncelli, E.P., 2003. Vision and the Statistics of the Visual Environment. *Current Opinion in Neurobiology* 13, 144-149.
68. Srinivasan M.V., Laughlin S.B. and A. Dubs, 1982. Predictive Coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society B* , 216, 427-459.
69. Stevens C.F., 2003. Neurotransmitter Release at Central Synapses. *Neuron* 40, 381-388.
70. Sugrue L.P., G.S. Corrado and W.T. Newsome, 2004. Matching behavior and the representation of value in the parietal cortex, *Science* 304, 1782-1787.
71. Sugrue L.P., G.S. Corrado and W.T. Newsome, 2005. Choosing the greater of two goods: neural currencies for valuation and decision making. *Nature Reviews Neuroscience* 6, 363-375.
72. Thompson R.F., 2000. *The Brain: A Neuroscience Primer*. 3rd Edition, Worth Publishers.
73. Van Essen D. C., C.H. Anderson and D.J. Felleman, 1992. Information processing in the primate visual system: an integrated systems perspective. *Science* 255, 419-423.

74. Van Zandt T., H. Colonius, R.W. Proctor, 2000. A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin and Review*, 7, 208-56.