



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

From perception to action: An economic model of brain processes[☆]Isabelle Brocas^{a,b}, Juan D. Carrillo^{a,b,*}^a Department of Economics, University of Southern California, Los Angeles, CA 90089, United States^b CEPR, United Kingdom

ARTICLE INFO

Article history:

Received 3 October 2010

Available online 20 October 2011

JEL classification:

D81

D83

D87

Keywords:

Neuroeconomic theory

Information processing

Bayesian updating

ABSTRACT

We build on research from neurobiology to model the process through which the brain maps outside evidence into decisions. The sensory system encodes information through cell-firing. Cell-firing is measured against a threshold, and an action is triggered depending on whether the threshold is surpassed. The decision system modulates the threshold. We show that the (constrained) optimal threshold is set in a way that existing beliefs are likely to be confirmed. We then derive behavioral implications. Our mechanism can explain in a unified framework a number of ‘anomalies’ noted in psychology and economics: (i) belief anchoring (the order in which evidence is received affects beliefs and choices); (ii) polarization (individuals with opposite priors may polarize their opinions after receiving identical evidence); (iii) payoff-dependence of beliefs and (iv) belief disagreement (individuals with identical priors who receive the same evidence may end up with different posterior beliefs).

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

Economic theory has traditionally been interested in the analysis of *choices*. In particular, the *processes* by which individuals reach decisions have been ignored, mainly because we had little knowledge of the pathways going from perception to action. With the recent developments in techniques to measure brain activity, the neurobiology and neuroscience literatures have substantially improved our understanding of the *biological* mechanisms that transform sensory perceptions into voluntary actions. These results can now be incorporated into formal economic models of decision-making.

The objective of this paper is to offer a brain-based model of information processing that builds on evidence from neurobiology. According to this evidence, there are five basic principles in the physiological mechanisms of information processing.¹ *First*, neurons carry information away from the sensory system using an imperfect encoding technology: the level of neuronal cell firing depends stochastically on the information obtained. *Second*, an action is triggered when the cell firing activity in favor of one alternative reaches a certain threshold. *Third*, neuronal activity responds to changes in payoffs and beliefs, that is, neurons compute approximately the ‘expected value’ associated to each alternative. *Fourth*, neurons

[☆] This paper is a major revision of the paper entitled “Reason, Emotion and Information Processing in the Brain”. We thank C. Camerer, E. Dal Bo, A. Daughety, E. Fehr, J.J. Ganuza, D. Houser, A. Rangel, J. Reinganum, J. Schall, S. Scotchmer, R. Seamans, G. Tan and seminar participants at U. Carlos III, the Behavioral Conference at Fundación Areces, USC, U. of Guelph, UC Davis, UC San Diego, ASSA 2008, the Conference on Neuroeconomics and the Law at USC, the Neuroeconomics Conference at the Arne Ryde Foundation, the CSLP Symposium on Cognitive-Neuroscience and the Law at USC, UC Berkeley Haas Business School, ITAM, Vanderbilt University and University of Edmonton for helpful comments.

* Corresponding author at: Department of Economics, University of Southern California, Los Angeles, CA 90089, United States.

E-mail addresses: brocas@usc.edu (I. Brocas), juandc@usc.edu (J.D. Carrillo).

¹ For brevity concerns, the evidence is summarily reviewed in Section 2 as we present each ingredient of the model. For the economist interested in a more detailed account of the relevant neurobiological literature, we refer to the survey by Brocas (2009).

also perform approximately 'Bayesian' inferences conditional on the data retained. *Fifth* and crucial for our purpose, the triggering threshold can be modified, and this affects how the evidence is interpreted.

Building on these premises, our analysis has three goals. First, we propose the first economic model capable of predicting decisions when the underlying mechanism that processes information exhibits the basic premises of neurobiology. Second, we characterize the decisions that result from this mechanism in environments with varying degrees of complexity. Third and most importantly, we analyze the behavioral implications for decision-making. In particular, we discuss behaviors that could not arise in a traditional learning model. We now review these three contributions.

In Section 2, we represent the underlying mechanism that maps sensory perception into actions. We start with a very simple, two-states (A and B), two-actions (a and b) model, where payoffs depend on the combination of action and state. The individual receives a signal from the outside world which is encoded by the sensory system in the form of cell firing. The information stochastically depends on the state, with high cell firing being more likely in state A and low cell firing in state B . Cell firing is measured against a threshold. If it is surpassed, action a is undertaken, otherwise action b is implemented. This model represents the underlying economical mechanism that transforms perception into action. According to the previous evidence, the threshold can be modified. We assume that the system (broadly defined) that transforms perception into actions – the decision system – selects the threshold that maximizes the expected payoff of the individual.² This representation is extended to complex and behaviorally more relevant environments that include many actions and states, and allow the decision system to set several thresholds simultaneously or sequentially. The model thus provides a common language for economists and neuroscientists interested in the study of decision processes.

In Section 3, we characterize information processing and decision-making. In the basic setting with two states and two relevant actions, the threshold is optimally set in a way that beliefs are likely to be supported. That is, if the agent becomes more confident that the state is A , the threshold is decreased. Thus, the new threshold is more likely to be surpassed whether the state is indeed A or not and, as a result, the agent is more likely to take the action which is optimal in that state but suboptimal in the other (Proposition 1). The logic for this property is simple: as the likelihood of A increases, stronger contradictory information is required to reverse that belief. This result is reminiscent of the findings obtained in the classical theory of organizations literature (Calvert, 1985; Sah and Stiglitz, 1986; Meyer, 1991) using related formal models. The contribution of the theory section is to extend the model to behaviorally relevant environments, where decisions are typically complex. We show that the result mentioned above holds under some technical conditions when there is a continuum of relevant actions (Proposition 2). It also extends to a dynamic setting with threshold re-optimization between stages (Proposition 3). In fact, the ability to modify neuronal thresholds has a snowball effect on decision-making: a stronger belief towards one state implies a greater threshold modulation in its favor, therefore a higher probability that new information supports it, and so on.

In Section 4, we provide some behavioral implications of this theory. This is the most critical contribution of the paper. Information processing in the brain has two major ingredients: the information received is noisy, and the decision threshold mechanism filters out some of it. In other words, the evidence received from the sensory system is *interpreted*. We analyze the effect of the interpretative feature of the mechanism, and focus on behaviors that would not emerge in a standard learning environment where the exact signals were processed. We obtain five implications. First, we show that the sequence in which signals are received affects posterior beliefs and decisions (Implication 1). Indeed, given the threshold mechanism induces the individual to confirm his prior beliefs, a first piece of information may act as an anchor for subsequent learning experiences. Second, individuals with different prior beliefs who receive identical signals may move their beliefs farther apart (Implication 2). This is again due to the interpretative feature of the threshold mechanism. Suppose two individuals are willing to take two different actions following their prior beliefs, releasing mixed evidence will reinforce their priors and individuals will polarize both their beliefs and their actions. Third, individuals with identical priors but different payoffs will hold different posterior beliefs even if they are exposed to the same evidence (Implication 3). Indeed, preferences shape beliefs: individuals who feel differently about the losses incurred when the wrong action is undertaken will set different thresholds and therefore interpret identical information in different ways. Fourth, the optimal decision-threshold mechanism generates payoff-dependent posterior beliefs (Implication 4). As such, the preferences of an individual are best represented by an expected utility function where probabilities are payoff-dependent. The formulation is thus reminiscent of the rank-dependent expected utility and the security-potential/aspiration models. Fifth, cognitive limitations that result in setting an insufficient number of thresholds generates elimination strategies (Implication 5). More precisely, the decision system disregards the most unlikely states and discriminates optimally among the remaining ones. Each of these five implications has received strong empirical and experimental support in social psychology and behavioral economics. In particular, an abundant literature on confirmatory biases documents anchoring and polarization effects consistent with our two first implications. The relationship between preferences and beliefs has also been extensively studied. Therefore, our model provides a representation of the underlying neuronal mechanisms that generate some well-known biases in a unified framework.³

² Maximization is the only assumption of the model which is not based on neurobiology evidence. We will not defend it on the basis of "loose" evolutionary considerations. Instead, we argue that the point of the paper is precisely to determine what type of decision-making we obtain in the best case scenario of *optimal* adjustment. For rigorous evolutionary models of human economic characteristics, see e.g. Robson (2001a, 2001b).

³ Our model is consistent with the "satisficing" model in Simon (1956), advocating that cognitive limitations should be taken into account to propose a more realistic approach to rationality. In our model, the decision-maker optimizes under physiological constraints.

Our work departs substantially from the existing neuroeconomics literature. Current theories are largely interested in conflicts between brain systems. Behavioral anomalies are the result of an interplay between a rational and an automatic system.⁴ Instead, we propose a model that incorporates the *physiological constraints faced by the brain in the process of decision-making*. Since we focus on the physiological mechanisms behind choice processes, our paper is closer in essence to the “physiological expected utility” theory developed by Glimcher et al. (2005).⁵ There are also several areas outside economics that study a related problem, although from a rather different angle. Neurobiologists have worked on complex statistical algorithms that mimic what neurons in the brain do (see e.g. Brown et al., 1998). Theoretical neuroscientists have constructed computational models of the brain based on the underlying biological mechanisms (see Dayan and Abbott, 2005, for an introduction). Psychophysicists have developed a “signal-detection theory” to study the likelihood of finding a weak signal in a noisy environment (see Wickens, 2002, for an introduction). Finally, there is a literature on neural networks and artificial intelligence which builds models inspired by the architecture of the brain in order to solve specific tasks like data processing or filtering (see Abdi, 1994, for an overview).

The paper is organized as follows. In Section 2, we present our model and review the neurobiology literature. In Section 3, we characterize the optimal thresholds and subsequent decisions in environments with varying levels of complexity. In Section 4, we develop several implications for choice under uncertainty. In Section 5, we provide some concluding remarks. All the proofs and some supplementary material (extensions, robustness of the theory, and analytical examples) are relegated to Appendix A.

2. Modeling brain processes

Consider a primitive individual (he) whose options are to stay in the cave or go hunting. On a dangerous day, the decision to go hunting can result in an injury. On a safe day, he may catch a prey and save his family from starvation. Before making the decision, the individual can take a look outside the cave. This conveys information about the hypothesis he should endorse: is it dangerous to leave the cave or is it safe?

The brain uses a specific mechanism to analyze situations and make decisions. Neurons carry information from the sensory circuitry, where information about the outside world is received, to the decision-making circuitry, where the information is aggregated and interpreted. The objective of this section is to provide a *parsimonious* model able to represent the brain mechanisms underlying the decision process. We will first concentrate on a two-action discrimination task. This is the situation that has been documented most extensively in neurobiology. To simplify the exposition, we will use the previous example as an illustration. The model is generalized afterwards.

We build our theory around two well-known paradigms, both related to decision-making under risk or uncertainty, where subjects choose an option that yields a stochastic payoff. One paradigm focuses on *utility evaluation*. A monkey is offered alternatives with different payoffs and different probabilities, and must pick one. Probabilities and rewards are chosen by the experimenter and transmitted as objectively as possible to the subject (this is the setting pioneered by Platt and Glimcher, 1999). Typically, the monkey faces the following problem: there are two alternatives, $i \in \{a, b\}$, yielding payoff π_i with probability q_i and 0 with probability $1 - q_i$. An objective of this line of experiments is to vary the magnitude of payoffs π_i and probabilities q_i and correlate them with the activity of neurons believed to transform visual data into commands. The goal is to determine whether neuronal cell firing is proportional to the expected value of the reward. The other paradigm focuses on *belief updating*. A monkey is offered two options and must infer from a noisy signal which one is objectively correct (this line of work uses variants of the random dot stimuli experiment, see e.g. Newsome et al., 1989 and Salzman et al., 1990). In this experiment, the monkey must form a belief, that is, he must determine the probabilities he should apply to the choice problem. Formally, there are again two alternatives, $i \in \{a, b\}$, each yielding payoff π_i or 0 but now the probabilities are not transmitted to the subject. Instead, they must be inferred and computed. Experiments in this literature vary the noise to signal ratio and correlate the activity of neurons tuned to detect features of the signal with choice. The goal is to determine whether neurons update information in an approximately Bayesian way.

Our decision-making problem combines aspects of both paradigms. As in the first paradigm, we relate choices to posterior beliefs and magnitude of rewards. As in the second paradigm, we derive posterior beliefs from prior knowledge and noisy information. In the following subsections, we build our model step by step.

2.1. Building blocks of the theory

There are three ingredients necessary to make a decision.

2.1.1. Environment

It consists in the set of alternatives and states, which are exogenously given. In our example, the alternatives are two actions, a (stay in the cave) and b (go hunting) and the states are the number of predators out there, A (dangerous day

⁴ See Thaler and Shefrin (1981) and Shefrin and Thaler (1988) for the seminal contributions, and Zak (2004) or Camerer et al. (2005) for reviews.

⁵ There is also a recent experimental literature that explores the neurobiological foundations for social behavior (see e.g. Zak et al., 2004, on the effect of oxytocin on trust, Spitzer et al., 2007, for a study of the neural mechanism underlying social compliance and Fehr and Camerer, 2007, for an overview).

with many predators) and B (safe day with few predators). The action space is denoted by Γ with typical element γ , and the state space is denoted by S with typical element s .

2.1.2. Preferences over outcomes

They reflect how the individual feels about the consequences of his decision. These also correspond to the monkey's assessment of the rewards in the neurobiology experiments. In our example, there is an optimal action in each state: it is better to stay in the cave on a dangerous day and to go hunting on a safe day. When actions and states are not matched properly, the individual incurs losses (injury, starvation). This can be represented by a loss function $\tilde{l}(\gamma, s)$ where $\gamma \in \{a, b\}$ and $s \in \{A, B\}$, $\tilde{l}(a, A) = \tilde{l}(b, B) = 0$ and $\tilde{l}(\gamma, s) < 0$ otherwise. The individual must form a representation of the losses he might incur. These losses can be objective or subjective (e.g., individuals may differ in valuing the cost of an injury, or the benefit of food for the family). We will denote by \mathcal{L} the complete description of outcomes. It is a map from $\Gamma \times S$ into \mathbb{R}_- . In our example, it is simply $\mathcal{L} = \{\tilde{l}(a|A), \tilde{l}(b|A), \tilde{l}(a|B), \tilde{l}(b|B)\}$.

2.1.3. Information

The individual makes inferences from the information available. There are two types of knowledge. First, before taking a look out of the cave, the individual has a sense of how likely each state is. He may use memories of previous episodes to make that assessment. This information can be summarized in a *prior belief probability* $p \in (0, 1)$ that the state is A . Despite its importance, we will ignore how priors are formed and simply assume that the individual is able to make a decision if he does not get to look outside the cave. Second, the individual receives a *signal* about the outside world which is collected by the sensory circuitry (visual and auditory cortices for instance) and processed to compute a posterior before making a choice. We now discuss what a signal in the brain is.

2.2. Signals in the sensory system

The information received from the outside world is encoded and translated into neuronal activity. This activity represents the signal in the brain. We will capture it with the continuous variable $c \in [0, 1]$, ordered from safest ($c = 0$) to most dangerous ($c = 1$) perceived environment. The mechanism underlying the construction of c can be described as follows. Each neuron detects 'danger' or 'no danger'. The variable c can be interpreted as the fraction of neurons that detect 'danger', and therefore favor the hypothesis that the state is A .⁶ This formalization encompasses two polar cases: (i) the image received from the outside world is fully informative but some neurons fire in the 'wrong' direction (mistakes in interpretation), and (ii) the image received is already distorted by external factors (noise in the environment).

The existence of an entire range of c captures the stochastic variability in neuronal cell firing. The evidence from neurobiology in favor of this hypothesis is overwhelming. Even when exposed to the same stimuli, neurons do not always fire in the same way. Different neurons detect different features and sometimes compete (Ma et al., 2006; Nichols and Newsome, 2002). Stochastic variability also captures the metabolic costs associated to activity (Laughlin et al., 1998), and the existence of noise in the process (e.g., stochastic neurotransmitter release Stevens, 2003). Part of the variability can also be due to the context in which the image is received (e.g., naturalistic conditions Simoncelli, 2003). Finally, initial signals are aggregated within and across populations of neurons that detect different features which, again, results in noisy aggregation (Shadlen et al., 1996).

At the same time, perception is intimately related to state. Formally, when the state is s , the likelihood of c is $f(c|s)$, with $F(c|s) = \int_0^c f(y|s) dy$ representing the probability of a cell firing activity not greater than c . To capture the idea that high cell firing (a majority of neurons carrying the signal danger) is more likely to occur when $s = A$ and low cell firing (a majority of neurons carrying the signal no danger) is more likely to occur when $s = B$, we impose a standard Monotone Likelihood Ratio Property (MLRP) assumption⁷:

Assumption 1 (MLRP).

$$\frac{\partial}{\partial c} \left(\frac{f(c|B)}{f(c|A)} \right) < 0 \quad \text{for all } c \quad (\text{A1})$$

In this paper, we do not discuss the origin of the function $f(\cdot|s)$. The particular functional form may be specific to the task, or specific to the individual, or to the individual for the task. The function may be shaped over time through repeated exposure. For our purpose, it is simply given.

⁶ Signals may also be received independently from different sensory systems indicating how likely a danger is. In that case, c can be seen as the aggregate of this information.

⁷ MLRP implies: (i) $\frac{1-F(c|A)}{f(c|A)} > \frac{1-F(c|B)}{f(c|B)}$, (ii) $\frac{F(c|A)}{f(c|A)} < \frac{F(c|B)}{f(c|B)}$, and (iii) $F(c|A) < F(c|B) \forall c \in (0, 1)$. Note that several papers in computational neurobiology adopt special cases of MLRP with specific functional forms. For example, Ma et al. (2006) argue that a Poisson noise fits the data well. Formally, let $d(s)$ be the tuning curve of a neuron (which peaks at the state s the neuron detects best). The density is then $f(c|s) = \frac{e^{-d(s)} d(s)^c}{c!}$. The best fit satisfies $d(B) < d(A)$, which is more restrictive than MLRP.

2.3. Decision mechanism in the decision system

A decision is made on the basis of the overall knowledge of the problem. This knowledge consists in preferences \mathcal{L} and information (the prior p and the signal c). A decision is then a map between knowledge and alternatives. Formally, it is a function \mathcal{D} that assigns an action in Γ to any triple (\mathcal{L}, p, c) . The objective of this section is to characterize this map under the assumption that (i) decision-making in the brain is an economical process and (ii) its representation must be consistent with actual neuronal cell-firing.

2.3.1. Decision-thresholds

Neuronal thresholds and synaptic connections filter information. Depending on how high neuronal thresholds and/or how strong synaptic connections are, neuronal activity will be stopped or propagated along a given path and trigger an action. This mechanism is economical in that it requires minimal knowledge to reach a decision. In a classical study, Hanes and Schall (1996) use single cell recording to analyze the neural processes responsible for the duration and variability of reaction times in monkeys. The authors find that movements are initiated when *neural activity reaches a certain threshold activation level*, in a winner-takes-all type of contest. This evidence suggests that the mechanism can be represented by a decision-threshold mechanism: it is *as if* there exists a threshold x such that action b is triggered when $c < x$, and action a is triggered when $c > x$. At the same time, it filters information out. In other words, the mechanism provides an ‘interpretation’ of the information. The sensory system collects c and the decision system interprets it as $c < x$ or $c \geq x$. The decision system compares alternatives via this mechanism (see Shadlen et al., 1996; Gold and Shadlen, 2001 and Ditterich et al., 2003, for further evidence).⁸ Overall, decision-making can be represented by a decision-threshold mechanism of the form⁹:

$$\mathcal{D}(\mathcal{L}, p, c) = \begin{cases} a & \text{if } c \geq x \\ b & \text{if } c < x \end{cases}$$

2.3.2. Expected utility evaluation

As briefly introduced in Section 2, one strand of the neuroscience literature is interested in testing for ‘argmax-like’ decision making. The basic question is: can we represent the objective of the decision system as an expected payoff function? The answer is yes. Indeed, conditional on a given structure of payoffs and probabilities, this literature studies how alternatives are compared and selected and how this correlates with brain activity. A series of single neuron recording experiments with primates have demonstrated that both changes in beliefs and changes in payoff magnitudes are correlated with neuronal activity in the lateral intraparietal area (Platt and Glimcher, 1999; Padoa-Schioppa and Assad, 2006, 2008). Hence, those neurons carry information regarding both the probabilities to apply to the decision (in our case, the posterior beliefs) and the ensuing rewards, and compute approximately the “expected value” associated to each alternative (see also Roitman and Shadlen, 2002 and Glimcher et al., 2005). For our model it implies that, conditional on having reached a given posterior belief, the decision system will evaluate each alternative in a way compatible with expected utility theory. Formally, when $c < x$ taking the recommended action b is valued at¹⁰:

$$L(\Pr(A|c < x), b) = \Pr(A|c < x)\tilde{I}(b|A) + (1 - \Pr(A|c < x))\tilde{I}(b|B)$$

and when $c \geq x$ taking the recommended action a is valued at:

$$L(\Pr(A|c \geq x), a) = \Pr(A|c \geq x)\tilde{I}(a|A) + (1 - \Pr(A|c \geq x))\tilde{I}(a|B)$$

2.3.3. Bayesian inferences

Another strand of the literature tests for belief updating. The question here is: can we represent the behavior of neurons as performing approximately Bayesian inferences, that is, as computing posterior probabilities conditional on the data retained? The answer again is yes. One of the early theories, the “Efficient Coding Hypothesis” postulates that neurons encode information as compactly as possible, so as to use resources efficiently (Barlow, 2001; Simoncelli, 2003). This theory has recently led to a myriad of sophisticated statistical models that describe Bayesian stochastic processing of information by neurons in visual, auditory and haptic perception tasks (see e.g. Schwartz and Simoncelli, 2001; Ernst and Banks, 2002; Körding and Wolpert, 2004 and Ma et al., 2006). Building on the work by Hanes and Schall (1996), Shadlen et al. (1996) and Gold and Shadlen (2001) study a motion discrimination task, where monkeys must decide whether the net direction of dots that appear on a monitor is upward or downward. The authors develop a theory of how information is processed. It is shown that neurons “compute” approximately the log-likelihood ratio of the alternatives in order to determine which hypothesis should be supported by the evidence. Thus, according to this result, neurons incorporate the two major ingredients of Bayesian theory: prior probabilities, and stochastic information processing (see also Deneve et al., 1999, for a numerical

⁸ Similar models are used in other related settings. For instance, “explore vs. exploit” models of brain activity also adopt a decision threshold representation (see McKenzie et al., 2009, for a discussion).

⁹ Some models in the neuroscience literature are closely related to ours (e.g. Gold and Shadlen, 2001) while others are also interested in how information is accumulated through time. We discussed those in Section 2.6.

¹⁰ It is important to realize that the role of the probabilities in Platt and Glimcher (1999) are played by our posterior beliefs, not the priors.

simulation model). For our model, it implies that the likelihood that the state is A when $c > x$, can be estimated at its Bayesian posterior:

$$\Pr(A|c \geq x) \equiv \bar{p}(x) = \frac{[1 - F(x|A)]p}{[1 - F(x|A)]p + [1 - F(x|B)](1 - p)} \quad (1)$$

Similarly, if $c < x$, the likelihood that the true state is A is:

$$\Pr(A|c < x) \equiv \underline{p}(x) = \frac{F(x|A)p}{F(x|A)p + F(x|B)(1 - p)} \quad (2)$$

A trivial implication of Bayesian updating is that, for any p and x , the belief about state A is revised upwards if x is surpassed ($\bar{p}(x) > p$) and downwards if x is not reached ($\underline{p}(x) < p$). This captures the idea that low cell firing is an imperfect indicator of state B and high cell firing is an imperfect indicator of state A . Also, for a given p , suppose that x is surpassed. Then, the posterior belief about state A is revised upwards more strongly if c is more tightly correlated with the true state (in a stochastic dominance sense). This is consistent with evidence obtained in neurobiology: using a similar experiment as Shadlen et al. (1996), Ditterich et al. (2003) show that when the task is more difficult (fewer dots move in synchrony), monkeys make more mistakes.¹¹

Notice that an outside observer may conclude the subject holds non-Bayesian beliefs when those are elicited. This is the case because updating is done through an internally consistent but constrained Bayesian process, where only the summary statistic $c < x$ or $c \geq x$ is retained.¹² So, for example, two different signals one stronger than the other but both above the threshold x will be interpreted identically. In Section 4 we extensively use this property of our model to discuss behavioral implications that depart from standard theories of learning.

2.4. Optimal modulation

As explained before, the decision-threshold represents the mechanism through which information from the sensory system is translated into decisions. There is ample evidence that neuronal thresholds and synaptic connectivities are modulated, modifying the conditions under which information is propagated along a pathway. The effect of setting a threshold is to decide how to interpret the evidence. At the level of neurons, this can be achieved by a combination of tools. First, neurotransmitter releases can affect neuronal thresholds (see Mogenson, 1987). Neurons that carry signals of ‘danger’ will fire only if they receive a sufficiently strong cue regarding predators. If their threshold is increased, they are less likely to fire. This also means that neurons will pass over a ‘danger’ signal more infrequently. Second, synapses filter the information that is just relevant to make a decision (by adjusting their strength, see Klyachko and Stevens, 2006). In terms of our model, this means that x can and will be adjusted. Intuitively, it is natural that a high level of cell-firing should be interpreted differently depending on the preferences and prior belief of the individual. As such, modulation must be the mechanism used by the decision system to account for \mathcal{L} and p .¹³

Given the possibility of modulation, we want to construct a mechanism that sets the threshold optimally. As discussed in footnote 2, optimality is an exploratory assumption, which allows us not to impose any ad-hoc constraint on decision-making. It requires two conditions. First, x must be such that the optimal action is different depending on whether the threshold is reached or not. Formally, $L(p(x), b) \geq L(p(x), a)$ and $L(\bar{p}(x), b) \leq L(\bar{p}(x), a)$. However, many thresholds satisfy this property. Let $X(\mathcal{L}, p)$ be the set of all such thresholds. At equilibrium, the optimal threshold $x^*(\mathcal{L}, p)$ must be in that set.¹⁴ Second, among those, the optimal threshold is such that the decision rule yields highest expected utility. Formally, it solves:

$$\arg \max_{x \in X} V(x) = \Pr(c > x)L(\bar{p}(x), a) + \Pr(c < x)L(\underline{p}(x), b)$$

Overall, if modulation serves the purpose of improving decision making, the threshold will be set in such a way that these two conditions are satisfied.¹⁵

¹¹ Similar results have also been obtained with human subjects using fMRI studies. For instance, Heekeren et al. (2004) find that the mechanism by which the brain of a monkey computes perceptual decisions is also at work for humans and for more sophisticated choices, such as image recognition.

¹² From a technical viewpoint, our model with coarse information partition is related to the recent literature on categorical thinking (Fryer and Jackson, 2008), where individuals lump experiences in a limited set of categories, and this categorization subsequently affects their decisions. An advantage of our approach is that it is based on the physiology of information processing.

¹³ Interestingly, even though it is well known that thresholds can be modulated, the existing literature considers models with fixed thresholds and does not optimize over x . This is arguably the major difference between our constrained optimization approach and all the literature in computational neurobiology.

¹⁴ When $x \notin X(\mathcal{L}, p)$, the subject takes the same action independently of whether the threshold is reached or not. It means that information is irrelevant for decision-making. This is necessarily dominated by an informative signal.

¹⁵ The mechanism could be implemented via a Hebbian type of learning mechanism (Hebb, 1961), which postulates that an increase in synaptic strength arises from the presynaptic cell's repeated and persistent stimulation of the postsynaptic cell. The activation patterns in one episode (e.g., hunting in the past) should therefore affect the way information is interpreted in the next one (looking outside the cave in the present).

2.5. General representation

The previous representation can be easily extended to include arbitrarily many states and actions. Suppose the decision system must discriminate between n alternatives from Γ , each indexed by i , and denote by $\mathcal{P}(S)$ the prior probability distribution over states. The decision process can be represented by a mechanism with $n - 1$ thresholds:

$$\mathcal{D}(\mathcal{L}, \mathcal{P}(S), c) = \begin{cases} 1 & \text{if } c \in [0, x_1) = A_1 \\ 2 & \text{if } c \in [x_1, x_2) = A_2 \\ \dots & \\ n & \text{if } c \in [x_{n-1}, 1] = A_n \end{cases}$$

Alternative i 's value is $L(\mathcal{P}(S|c \in A_i), i) = E_S[\tilde{I}(i|s)|c \in A_i]$. Let $X_i(\mathcal{L}, \mathcal{P}(S))$ be the set of thresholds x_i such that it is optimal to take action i between x_{i-1} and x_i and let $V(x_1, \dots, x_{n-1}) = E_{A_i}[L(\mathcal{P}(S|c \in A_i), i)]$. Finally, denote by $x_i^*(\mathcal{L}, \mathcal{P}(S))$ the optimal threshold. We obtain the following representation:

Representation. In a generic n -action discrimination task, the decision is represented by the following decision-threshold mechanism:

$$\mathcal{D}(\mathcal{L}, \mathcal{P}(S), c) = \begin{cases} 1 & \text{if } c \in [0, x_1^*(\mathcal{L}, \mathcal{P}(S))) \\ 2 & \text{if } c \in [x_1^*(\mathcal{L}, \mathcal{P}(S)), x_2^*(\mathcal{L}, \mathcal{P}(S))) \\ \dots & \\ n & \text{if } c \in [x_{n-1}^*(\mathcal{L}, \mathcal{P}(S)), 1] \end{cases}$$

where

- (i) $x_i^*(\mathcal{L}, \mathcal{P}(S)) \in X_i(\mathcal{L}, \mathcal{P}(S))$ for all i ;
- (ii) $x_i^*(\mathcal{L}, \mathcal{P}(S)) = \arg \max_{x_i \in X_i(\mathcal{L}, \mathcal{P}(S))} V(x_1, \dots, x_{n-1})$ for all i .

Naturally, as the number of alternatives grow, it will be more difficult for the brain to implement the optimal mechanism.

2.6. The diffusion model

At this stage, it is instructive to compare our threshold model with the “diffusion model”, a leading approach in neuroscience (see e.g. the survey by Ratcliff and McKoon, 2008). Although it has not been formalized exactly in those terms, the diffusion model is close to the two-states ($S \in \{A, B\}$), two-actions ($\gamma \in \{a, b\}$) “investment under uncertainty” model (Dixit and Pindyck, 1994), where action a is optimal in state A and action b is optimal in state B .¹⁶ The probability of A is initially assessed at $p \in (0, 1)$. Finally, information about the likelihood of A is accumulated dynamically and follows a Wiener process with a drift that depends on the state:

$$dp = v(S)dt + \sigma \varepsilon_t \sqrt{dt} \quad \text{where } v(S) = \begin{cases} m & \text{if } S = A \\ -m & \text{if } S = B \end{cases}$$

with $m > 0$. In this dynamic setting, as long as the belief p is in a certain interval (p^*, p^{**}), the individual keeps accumulating evidence. If the belief reaches the upper bound p^{**} he takes action a , whereas if it reaches the lower bound p^* he takes action b .

The diffusion model shares important similarities with ours. First, there is a stochastic relation between the state and the information perceived (the cell firing function $f(\cdot|S)$ and the drift $v(S)$). Mistakes occur in equilibrium and their frequency depends, among other things, on the accuracy of the aforementioned relation. Second, perception is continuous (cell firing c and information accumulation dp) but there are only two terminal posterior beliefs: $\Pr(A|c > x)$ and p^{**} trigger action a whereas $\Pr(A|c < x)$ and p^* trigger action b .

The diffusion model has one significant advantage over ours. Because it formalizes the dynamic information accumulation process, it has a rich set of predictions regarding *reaction time* (i.e., the lag between the moment the task is presented and the moment the choice is made) and its relationship with precision of information and likelihood of mistakes. This feature is not present in our model. On the other hand, our model also has some advantages: it can be easily extended to three or more actions and three or more states, and it is tractable enough to address formally optimal threshold modulation. This second point is key since the entire premise of our analysis (and of the neuroscience evidence) rests on the ability of the

¹⁶ For example, in the typical random dot stimuli experiment, the state is the direction of the majority of dots, upwards or downwards, and the action is the subject's best guess, upwards or downwards.

individual to modify thresholds and our goal is to study the implications of such modulation for decision-making.¹⁷ Overall, we believe that the model we build is the most appropriate given both the evidence and our objectives.

Finally, notice that our model is identical to a diffusion model with optimal bounds, no cost of information and a fixed deadline for making a decision. Indeed, in the absence of a cost, it is optimal to exhaust the accumulation of information and choose *A* or *B* depending on the final posterior. In that formulation, *c* corresponds to the integral of the information accumulated.¹⁸

3. Optimal thresholds and action selection

In this section we characterize optimal thresholds and choices under different assumptions about the number of alternatives and states. To isolate the effect of each assumption, we study each case separately.

3.1. Choosing between many actions

Suppose there are only two states $S = \{A, B\}$ and a continuum of actions $\Gamma = [0, 1]$, where lower values of γ denote going farther away from the cave to hunt. The preferences over outcomes are captured by the following loss functions:

$$\tilde{l}(\gamma, A) = \pi_A l(\gamma - 1), \quad \tilde{l}(\gamma, B) = \pi_B l(\gamma - 0)$$

where $l(z) = l(-z)$ for all z , $l'(z) < 0$ for all $z > 0$, and $\pi_s > 0$. The individual should go far away on safe days ($\gamma = 0$ if $s = B$) and stay close on dangerous days ($\gamma = 1$ if $s = A$).¹⁹ Note that the marginal cost of taking a wrong action in state s is proportional to π_s . Therefore, $\pi_A > \pi_B$ means that hunting on a dangerous day is more costly than staying in the cave on a safe day (e.g., the probability of a fatal injury may be greater than the probability of starvation). Given a posterior belief μ , the expected payoff of taking action γ is:

$$L(\mu, \gamma) = \mu[\pi_A l(\gamma - 1)] + (1 - \mu)[\pi_B l(\gamma)] \tag{3}$$

Suppose that $l(z)$ is weakly convex on both sides of its bliss point $z = 0$. Formally, $l''(z) \geq 0$ for all $z \neq 0$ so that departures from the optimal action are decreasingly costly.²⁰ In that case, $L(\mu, \gamma)$ is weakly convex in γ and differentiable on $(0, 1)$, so corner solutions are optimal. Denoting by $\gamma^*(\mu) = \arg \max_{\gamma} L(\mu, \gamma)$ when $l(z)$ is weakly convex, using (3) and given $l(0) > l(1)$, we have:

$$\gamma^*(\mu) = \begin{cases} 1 & \text{if } L(\mu; 1) \geq L(\mu; 0) \Leftrightarrow \mu \geq p^* \equiv \pi_B / (\pi_A + \pi_B) \\ 0 & \text{if } L(\mu; 1) < L(\mu; 0) \Leftrightarrow \mu < p^* \equiv \pi_B / (\pi_A + \pi_B) \end{cases}$$

Notice that $dp^*/d\pi_A < 0$ and $dp^*/d\pi_B > 0$: if the marginal cost of an incorrect action in a given state increases, then the individual is more willing to take the action optimal in that state even at the increased risk of erring in the other state. In our example, as predators become smarter and more dangerous, the individual is more likely to decide to stay in the cave, even on days that are apparently safe.

Suppose now that $l(z)$ is strictly concave. Formally, $l''(z) < 0$ for all z so that departures from the optimal action are increasingly costly.²¹ Denote by $\gamma^{**}(\mu) = \arg \max_{\gamma} L(\mu, \gamma)$ when $l(z)$ is concave. Taking the first-order condition in (3), we have:

$$\frac{\pi_B}{\pi_A} \frac{l'(\gamma^{**}(\mu))}{l'(1 - \gamma^{**}(\mu))} = \frac{\mu}{1 - \mu} \tag{4}$$

where $\frac{\partial L(\mu, \gamma)}{\partial \gamma} |_{\gamma=0} = -\pi_A \mu l'(1) > 0$, $\frac{\partial L(\mu, \gamma)}{\partial \gamma} |_{\gamma=1} = \pi_B (1 - \mu) l'(1) < 0$, and $\frac{\partial^2 L(\mu, \gamma)}{\partial \gamma^2} < 0$. In this case, the optimal choice is always interior.²² Lemma 1 summarizes these findings.

¹⁷ The bounds in the diffusion model where the individual stops accumulating evidence can, in principle, be optimized (as it is done in the investment under uncertainty literature). However, this requires a cost of obtaining (or waiting for) information, otherwise it is optimal to set the bounds at 0 and 1. Moreover, the location of these bounds will depend crucially on the magnitude of that cost. In our view it is hard to justify such a “cost” based on neural data. Some papers in neuroscience propose an optimization of the bounds (see Gold and Shadlen, 2002 and Bogacz, 2007, for a review), but the objective function that is maximized is not a discounted expected utility. It would be interesting to study the relationship between the optimal bounds derived in that literature and the optimal bounds emerging from a formal economic model based on expected utility.

¹⁸ Interestingly, this is quite close to the random dot stimuli setting, where monkeys first observe the movements of dots for a fixed amount of time, then they make their choice, and finally they obtain a reward.

¹⁹ Most of our results could be extended to the case of asymmetric loss functions $l_A(\cdot)$ and $l_B(\cdot)$. However, we already have two asymmetries in the model: $f(c|B)$ vs. $f(c|A)$ and π_A vs. π_B . Therefore, little would be gained in terms of generality by adding a third asymmetry. Technically, it would simply amount to replacing the expressions $\pi_B l(\gamma)$ with $\pi_B l_B(\gamma)$ and $\pi_A l(\gamma - 1)$ with $\pi_A l_A(\gamma - 1)$.

²⁰ This implicitly requires $l(z)$ to have a kink at $z = 0$. A special case would be the linear loss function $l(z) = -|z|$.

²¹ In this case, there may or may not be a kink at $z = 0$.

²² Under more general specifications of $l(\cdot)$, the optimal choice with a concave loss could be at the boundary ($\gamma^{**} \in [0, 1]$). What is important for our theory is that weak convexity or concavity of $l(\cdot)$ determines whether the solution has to be in a corner or not necessarily.

Lemma 1. Suppose $S = \{A, B\}$ and $\Gamma = [0, 1]$. The optimal action is generically unique and corner if $l''(z) \geq 0$ ($\gamma^*(\mu) \in \{0, 1\}$) and unique and interior if $l''(z) < 0$ ($\gamma^*(\mu) \in (0, 1)$).

We can now characterize the optimal threshold mechanism in both situations.

3.2. Optimal threshold modulation when the loss function is convex

Given Lemma 1, when $l''(z) \geq 0$ the problem boils down to a discrimination task with two states and two relevant actions. Applying our representation, the value function is:

$$\begin{aligned} V(x) &= \Pr(c > x)L(\bar{p}(x), 1) + \Pr(c < x)L(\underline{p}(x), 0) \\ &= p\pi_A[(1 - F(x|A))l(0) + F(x|A)l(1)] + (1 - p)\pi_B[(1 - F(x|B))l(1) + F(x|B)l(0)] \end{aligned} \tag{5}$$

The optimal threshold is $x^* = \arg \max_x V(x)$, which leads to the following result.

Proposition 1. When $l''(z) \geq 0$, the optimal process involves a single decision threshold x^* which is given by:

$$\frac{f(x^*|B)}{f(x^*|A)} = \frac{p}{1 - p} \frac{\pi_A}{\pi_B} \tag{6}$$

The optimal threshold implies that $dx^*/dp < 0$ and $dx^*/d\Pi < 0$ where $\Pi = \pi_A/\pi_B$.

Proof. See A.1. \square

The result has three implications. First and foremost, the threshold is set in such a way that existing beliefs are likely to be confirmed. To see this, consider a symmetric situation with a prior belief $p > 1/2$. Setting a high threshold is not efficient: whether it is surpassed or not, the individual will still think that A is the most likely state. Instead, setting a low threshold is optimal. If it is surpassed, the individual will slightly increase his confidence in state A . If it is not reached, he will become convinced that the state is B . The individual will end up taking different actions depending on the signal and, in both cases, he will be 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. In order to achieve this balance, the threshold should be low whenever A is a priori more probable than B and high otherwise.

Second, the optimal threshold depends on the relative payoffs. If the loss of taking the wrong action when $S = A$ increases, the threshold decreases, so it will be surpassed with higher probability. In equilibrium, the most costly mistakes are most likely to be avoided.²³

Third, the choice is efficient in the sense that the individual would not gain anything if he could observe the exact cell firing. Formally and using (6): $\Pr(A|c = x^*) = \pi_B/(\pi_A + \pi_B)$. This means that, $\gamma^* = 1$ dominates $\gamma^* = 0$ for all $c > x^*$ and $\gamma^* = 0$ dominates $\gamma^* = 1$ for all $c < x^*$. This, in turn, implies that for the purpose of choosing an action, it is sufficient to learn whether x^* is surpassed or not. The result hinges on the weak convexity of the loss function. Convexity reflects the fact that departures from the ideal choice are, at the margin, the most costly ones. Therefore, it is suitable to model environments where life-threatening events occur as soon as the optimal action is not taken.

It is interesting to note that our optimal threshold is in line with the proposal made in Gold and Shadlen (2001). Based on neurobiology evidence, the authors hypothesize that neurons form a decision variable that approximates the logarithm of the likelihood ratio, and the best decision under one hypothesis is triggered when this variable is above (or below) a decision threshold. They also hypothesize that the threshold affects how the evidence is interpreted and should be an indicator of the possible expected values of the outcomes. Eq. (6) formalizes this mechanism.

Finally, notice that Proposition 1 is reminiscent of Calvert (1985), Sah and Stiglitz (1986) and Meyer (1991), with the obvious differences in interpretations. The theoretical contribution of the paper consists in determining the optimal threshold(s) in more complex situations such as concave losses and dynamic choices.

3.3. Optimal threshold modulation when the loss function is concave

Suppose now that $l''(z) < 0$. As shown in Section 3.1, the optimal action can be anywhere in $(0, 1)$. An optimal process would involve a continuum of thresholds. Because this is difficult to implement in the brain, we determine what happens when only one decision-threshold x can be used. The solution is not first-best but it is a constrained maximum.²⁴ More

²³ Note that (A1) ensures uniqueness of maximum. The assumption, although accepted in neurobiology (see footnote 7), is not necessary for our main result. Indeed, if we adopted a substantially weaker first-order stochastic dominance assumption ($F(c|B) > F(c|A)$ for all c), then x^* would not necessarily be unique. However, since $V(x; p)$ is submodular, the comparative statics with respect to p and Π would be preserved in all the local maxima. The same submodularity argument can be applied to show that symmetry or continuity of $l(z)$ are not necessary either for the comparative statics to hold.

²⁴ We could extend the analysis to a situation with two or more thresholds. The resolution techniques would be similar. Naturally, the resulting inefficiencies would be reduced.

precisely, the brain sets x . If it is surpassed, the updated belief is $\bar{p}(x)$ and the action undertaken is $\gamma^{**}(\bar{p}(x))$. If it is not reached, the updated belief is $\underline{p}(x)$ and the action undertaken is $\gamma^{**}(\underline{p}(x))$. Given (1), (2) and (4), these actions satisfy:

$$\frac{\pi_B l'(\gamma^{**}(\bar{p}(x)))}{\pi_A l'(1 - \gamma^{**}(\bar{p}(x)))} = \frac{\bar{p}(x)}{1 - \bar{p}(x)} \left(= \frac{p}{1 - p} \frac{1 - F(x|A)}{1 - F(x|B)} \right) \tag{7}$$

$$\frac{\pi_B l'(\gamma^{**}(\underline{p}(x)))}{\pi_A l'(1 - \gamma^{**}(\underline{p}(x)))} = \frac{\underline{p}(x)}{1 - \underline{p}(x)} \left(= \frac{p}{1 - p} \frac{F(x|A)}{F(x|B)} \right) \tag{8}$$

where $\gamma^{**}(\bar{p}(x)) > \gamma^{**}(\underline{p}(x))$. Differentiating (7) and (8) and using (A1), we obtain:

$$\frac{d\gamma^{**}(\bar{p}(x))}{dx} > 0 \quad \text{and} \quad \frac{d\gamma^{**}(\underline{p}(x))}{dx} > 0$$

An increase in the threshold always induces the individual to choose a higher action. If it is surpassed, then the evidence in favor of A is stronger. If is not reached, then the evidence in favor of B is weaker. In both cases, higher actions follow. The new value function is identical to (5) except for the actions selected in either case. Formally:

$$V(x) = \Pr(c > x)L(\bar{p}(x), \gamma^{**}(\bar{p}(x))) + \Pr(c < x)L(\underline{p}(x), \gamma^{**}(\underline{p}(x))) \tag{9}$$

The optimal threshold discriminates optimally between the two actions. To find this threshold, we first need to introduce a strengthened version of MLRP.

Assumption 2 (*s-MLRP*). The probability distributions satisfy²⁵:

$$(i) \quad \frac{f(c|B)}{1 - F(c|B)} > \frac{f(c|A)}{1 - F(c|A)}, \quad (ii) \quad \frac{d}{dc} \left(\frac{f(c|B)}{f(c|A)} \frac{1 - F(c|A)}{1 - F(c|B)} \right) \leq 0, \quad (iii) \quad \frac{d}{dc} \left(\frac{f(c|B)}{f(c|A)} \frac{F(c|A)}{F(c|B)} \right) \leq 0 \tag{A2}$$

Define the following function:

$$H(x) \equiv F(x|B)l(\gamma^{**}(\underline{p}(x))) + (1 - F(x|B))l(\gamma^{**}(\bar{p}(x)))$$

We can now state our next result.

Proposition 2. *When $l''(z) < 0$, the optimal process requires a continuum of thresholds. A process that can discriminate only between two actions sets one threshold x^{**} and chooses actions $\gamma^{**}(\bar{p}(x^{**}))$ or $\gamma^{**}(\underline{p}(x^{**}))$. The optimal threshold satisfies:*

$$\frac{f(x^{**}|B)}{f(x^{**}|A)} = \frac{p}{1 - p} \frac{\pi_A}{\pi_B} \frac{l(1 - \gamma^{**}(\bar{p}(x^{**}))) - l(1 - \gamma^{**}(\underline{p}(x^{**})))}{l(\gamma^{**}(\underline{p}(x^{**}))) - l(\gamma^{**}(\bar{p}(x^{**})))} \tag{10}$$

It is unique and such that $dx^{**}/dp < 0$ if $dH(x^{**})/dx > 0$ and (A2) is satisfied.

Under (A1), $dH(x^{**})/dx > 0$ guarantees $dx^{**}/dp < 0$ in every locally optimal threshold but not uniqueness. Last, $dx^{**}/d\pi < 0$.

Proof. See A.2. □

When departures from the optimal action are increasingly costly, the quasi-concavity of the value function $V(x; p)$ is not guaranteed for generic values of $f(\cdot|A)$, $f(\cdot|B)$ and $l(\cdot)$. In fact, as x increases, two countervailing forces are at play. First and as before, the threshold is less likely to be surpassed and therefore more likely to induce the low action. Second, either outcome is a weaker indicator that the state is B . Therefore, the final action will be higher both when the threshold is surpassed and when it is not reached. Proposition 2 states that the main qualitative conclusions of Proposition 1 extend to a concave loss function as long as the problem is well-behaved, that is, if $dH(x^{**})/dx > 0$. The interpretation of this condition is simple: starting from the optimal threshold, a marginal increase in x increases the payoff of the individual if and only if the state is B . In other words, as x increases, the direct effect of increasing the likelihood of choosing the low action must dominate the indirect effect of choosing relatively higher actions. In Appendix A.2, we show that this condition is automatically satisfied when payoffs are quadratic ($l(z) = \alpha - \beta z^2$) and $\pi_A = \pi_B$. We also provide a complete characterization of the optimal threshold for that event, and an analytical solution of x^* and x^{**} in the linear and quadratic cases given specific functional forms for the distribution functions. As in Proposition 1, if the loss of taking the wrong action in a given state is increased, the individual will modify the threshold so as to favor that action.

Since concavity of $l(\cdot)$ captures a case where marginal departures from the ideal action are increasingly costly, this functional form is suitable to model environments where small mistakes are relatively harmless but large mistakes are very costly. Notice that not all the properties of the previous proposition hold. In particular, there is now an efficiency loss due to the inability of the individual to observe the exact cell firing level.²⁶ Last, it is important to realize that this process

²⁵ Note that (i) and (ii) or (i) and (iii) in (A2) imply (A1), but the converse is not true.

²⁶ This loss will be smaller the greater the number of thresholds.

does not prevent the individual from taking any action in Γ . Once a threshold is set, it discriminates only between two possible actions and only two possible posterior beliefs can be reached. However, x can be set to reach any action and any posterior belief from an ex-ante perspective. The result is then consistent with the experimental evidence according to which individuals can and often do report beliefs that vary across trials.

3.4. The dynamics of look-ups

In this section, we go back to the case of a convex loss function (hence, corner optima) and assume that the individual obtains two signals sequentially. The objective is to determine whether the tendency to confirm existing beliefs is mitigated or amplified in a more realistic setting where several pieces of information are processed sequentially. The analysis can also shed light on habituation, that is, how past signals affect the interpretation of current information. We impose the following characteristics on the representation of the dynamic process: (i) the process allows for threshold re-optimization between look-ups; (ii) it has memory, which means that the prior belief before the second look-up is simply the posterior belief after the first look-up; and (iii) it is forward-looking so that threshold modulation at each date takes into account future learning opportunities.

The rationale for these choices is the following. First, the dynamic setting adds a new dimension only under re-optimization between look-ups, otherwise it boils down to one look-up and a more accurate information. Property (i) guarantees that we are in this scenario. Second, we do not want to presuppose the existence of any exogenous loss of information between look-ups. Given property (ii), all the information collected and processed after a look-up is retained. Third, we want to concentrate on ‘intelligent processes’ that operate under constraints, rather than impose exogenous limitations on contingent planning. Property (iii) ensures the process optimizes at each period while anticipating future re-optimization. In other words, these choices put ourselves in a scenario which is interesting without being ad-hoc.

The dynamic choice can be modeled as the following two-stage problem. The individual initially holds a prior belief p . At stage 1, the sensory system collects c_1 . One or several thresholds are set, information is interpreted, and beliefs are updated. At stage 2, the sensory system collects c_2 . Again, one or several thresholds are set, information is interpreted and its recommendation is implemented. We assume that c_t is independently drawn from distribution $F_t(c_t|s)$ with $t \in \{1, 2\}$. Distributions may be different across stages but $\frac{f_t(c|B)}{f_t(c|A)}$ satisfies **(A1)** for all t .²⁷

Again, we concentrate on optimal modulation and extend our earlier representation. The choice problem can be represented by a sequence of decision-threshold mechanisms. Given (iii), those mechanisms obey backward induction. At date 2, we know that one of the corner solutions is optimal ($\gamma \in \{0, 1\}$), so one decision-threshold is enough (see Proposition 1). We denote the optimal threshold by x^* , since it is identical to that described in Section 3.2. At date 1, the process requires a continuum of thresholds (just like in the static case with a concave loss function described in Section 3.3). A constrained process that can discriminate only between two alternatives will set one forward looking decision-threshold that we denote by y^* . If it is surpassed, the posterior becomes $\bar{p}(y^*)$ and the optimal second stage threshold is $x^*(\bar{p}(y^*))$. If it is not reached, the posterior becomes $\underline{p}(y^*)$ and the optimal second stage threshold is $x^*(\underline{p}(y^*))$. Second stage thresholds are obtained using (1), (2) and (6). They satisfy:

$$\frac{f_2(x^*(\bar{p}(y))|B)}{f_2(x^*(\bar{p}(y))|A)} = \frac{\bar{p}(y) \pi_A}{1 - \bar{p}(y) \pi_B} \left(= \frac{p}{1-p} \frac{1 - F_1(y|A) \pi_A}{1 - F_1(y|B) \pi_B} \right) \tag{11}$$

$$\frac{f_2(x^*(\underline{p}(y))|B)}{f_2(x^*(\underline{p}(y))|A)} = \frac{\underline{p}(y) \pi_A}{1 - \underline{p}(y) \pi_B} \left(= \frac{p}{1-p} \frac{F_1(y|A) \pi_A}{F_1(y|B) \pi_B} \right) \tag{12}$$

The optimal threshold in the first stage maximizes the following value function:

$$W(y) = \Pr(c_1 > y)[V(x^*(\bar{p}(y)))] + \Pr(c_1 < y)[V(x^*(\underline{p}(y)))] \tag{13}$$

The first term is the likelihood of surpassing a cutoff y , in which case the posterior becomes $\bar{p}(y)$, multiplied by the second-stage value function given this posterior (see (5)), and under the anticipation of an optimal second-stage threshold $x^*(\bar{p}(y))$ (see (11)). The same logic applies to the second term. Notice that threshold y affects the utility of the individual only through its effect on the posterior belief. Define the following function:

$$J(y) \equiv F_1(y|B)F_2(x^*(\underline{p}(y))|B) + (1 - F_1(y|B))F_2(x^*(\bar{p}(y))|B)$$

We can now state our next result.

Proposition 3. *In a dynamic setting with $l'(z) \geq 0$, the optimal process requires a continuum of first stage thresholds. A process that can discriminate only between two alternatives in the first stage sets a threshold y^* that satisfies:*

²⁷ The setting assumes costless signals. It could be easily extended to costly signals, in which case the individual would also have to choose the optimal amount of information collected.

$$\frac{f_1(y^*|B)}{f_1(y^*|A)} = \frac{p}{1-p} \frac{\pi_A}{\pi_B} \frac{F_2(x^*(\underline{p}(y^*))|A) - F_2(x^*(\bar{p}(y^*))|A)}{F_2(x^*(\underline{p}(y^*))|B) - F_2(x^*(\bar{p}(y^*))|B)} \quad (14)$$

It is unique and such that $dy^*/dp < 0$ if $dJ(y^*)/dy > 0$ and **(A2)** is satisfied.

Under **(A1)**, $dJ(y^*)/dy > 0$ guarantees $dy^*/dp < 0$ in every locally optimal threshold but not uniqueness. Last $dy^*/d\pi < 0$.

Proof. See Appendix A.3. \square

Two-stage optimization problems are easily plagued by non-convexities in the overall maximand. Proposition 3 states that the qualitative conclusions of previous propositions are preserved in the dynamic version of the model if a technical condition, $dJ(y^*)/dy > 0$, is satisfied. As before, the intuition relies on the balance between the likelihood of the information and its impact. In fact, the two-stage model with decreasingly costly departures is technically similar to the one-stage model with increasingly costly departures. In particular, the same two effects operate when the threshold is increased. First, a direct effect: the new threshold is less likely to be surpassed. Second, an indirect effect: because surpassing a higher threshold is a stronger indicator of A and not reaching it is a weaker indicator of B , an increase in stage 1 threshold is always followed by a decrease in stage 2 threshold ($dx^*(\bar{p})/dy < 0$ and $dx^*(\underline{p})/dy < 0$). Condition $dJ(y^*)/dy > 0$ ensures that the direct effect dominates the indirect one. In Appendix A.3 we show that the condition automatically holds for any first-period distribution that satisfies **(A2)** if the second-period distributions are linear and symmetric.

Most importantly, the result highlights the dynamic snowball effect of threshold modulation on decision-making: a stronger belief towards one state implies a greater modulation in its favor, therefore a higher probability that new information supports it, and so on. We now provide below a simple analytical example that illustrates the theory.

Example 1. Suppose the density functions are identical in both stages, symmetric and linear: $f_t(c|A) = 2c$ and $f_t(c|B) = f_t(1-c|A) = 2(1-c)$, $t \in \{1, 2\}$. Let $\pi_s = 1$. From (6) and (14) and after some algebra, the optimal first and second stage thresholds are:

$$\left(\frac{1-y^*}{y^*}\right)^2 = \frac{p}{1-p} \Leftrightarrow y^*(p) = \frac{\sqrt{1-p}}{\sqrt{1-p} + \sqrt{p}} \quad \text{and} \quad \frac{1-x^*}{x^*} = \frac{p}{1-p} \Leftrightarrow x^*(p) = 1-p$$

Notice that $x^*(p) \geq y^*(p) \geq 1/2$ for all $p \leq 1/2$: for a given belief, the cutoff is always more extreme in the second than in the first stage. The intuition is that in the first stage the individual chooses the partition that conveys most information whereas in the second stage he chooses the partition that discriminates best among the two relevant alternatives. Finally, in Appendix A.3, we compute for this example the expected loss difference between using a process that employs a continuum of first stage thresholds and using a process that employs only one. For this particular example, the loss is small (2.5%).

A natural question would be to determine what happens as the number of information processing stages increases. Obviously, the final partition of beliefs becomes finer, which means that a greater number of posterior beliefs can be reached. Also, with three or more stages, the thresholds at all but the last stage only affect the belief inherited at the following stage. Thus, we conjecture that the main properties of the thresholds emphasized in Propositions 1 and 3 should, under reasonable conditions, be preserved (unfortunately, we have not been able to solve this problem analytically).

3.5. Summary

The results of Section 3 can be summarized as follows. As in the organizations literature, threshold processes are modulated in a way that existing beliefs are favored. Employing a one-threshold process (rather than a continuum of them) does not result in an efficiency loss in one-shot life-threatening situations (convex loss function). That result, however, does not extend to environments that are behaviorally more common such as the case of increasingly costly departures (concave loss function), and multiple look-ups. By contrast, the existence, qualitative properties of optimal threshold modulation, and comparative statics are all preserved in those environments under mild technical conditions. The theory extends in a number of other directions. In Appendix A.4, we show that all the comparative statics are also preserved if we consider a compact state space ($S \in [0, 1]$). In Appendix A.5, we show that optimal thresholds are more sensitive to initial beliefs in complex situations (where the correlation between cell firing and state is low) than in simple situations (where the correlation is high). Finally, the comparative statics obtained in this section and related appendixes are also in line with the results of single cell recording activity reported in neuroscience experiments. In the next section, we use these general comparative statics results to discuss implications for belief formation and choice under uncertainty.

4. Behavioral implications for choice under uncertainty

There are two main elements in our theory. First, the information encoded in the sensory system is stochastically correlated with the state. This is like a noisy report in a standard learning context. Second and most importantly, the threshold mechanism results in a loss of information. In particular, all signals on the same side of the threshold are pooled. This

means that the evidence received from the sensory system is incorporated in a Bayesian but coarse manner. In this section, we explore some implications of the model. We focus on behaviors that could not occur in a traditional environment where the exact signals were processed.

4.1. Confirmatory biases

In social psychology, a confirmatory bias is described as an error of inductive inference. It is the tendency of decision-makers to interpret evidence in a way that confirms their preconceived ideas about the world and avoid information that contradicts them (see Nickerson, 1998, for a review).²⁸ Our model offers a neurobiological mechanism for this bias. As developed in Section 3, when the belief that the state is A becomes stronger, the brain sets a lower threshold. Evidence is then more likely to be interpreted as endorsing A and less likely to be interpreted as endorsing B , both if the true state is A and if the true state is B . This simple principle has some subtle implications.

4.1.1. Belief anchoring and the role of first impressions

Suppose the decision-maker has a flat prior about an issue and receives a first piece of evidence that he uses to build a belief. According to our theory, once a belief is anchored, it is likely to be reinforced when the decision-maker is exposed to new information. This yields the following implication:

Implication 1 (Belief anchoring). *The sequence in which signals are received affects the beliefs and actions of the individual.*

Consider the dynamic setting of Section 3.4. Suppose for simplicity the environment is symmetric ($f(c|A) = f(1-c|B)$ and $\pi_A = \pi_B$) and the prior belief is $p = 1/2$, so that the first stage threshold is $y^*(1/2) = 1/2$. Consider the following symmetric signals \underline{c} and \bar{c} , with $\bar{c} = 1 - \underline{c}$, $\underline{c} \in (x^*(\Pr[A|c > 1/2]), 1/2)$ and therefore $\bar{c} \in (1/2, x^*(\Pr[A|c < 1/2]))$. If $c_1 = \bar{c}$ and $c_2 = \underline{c}$, the thresholds at both stages are surpassed. The decision-maker has a posterior belief $\Pr[A|c_1 > 1/2, c_2 > x^*(\Pr[A|c > 1/2])] > 1/2$ and takes action 1. Conversely, if $c_1 = \underline{c}$ and $c_2 = \bar{c}$, none of the thresholds is reached. The decision-maker has a posterior belief $\Pr[A|c_1 < 1/2, c_2 < x^*(\Pr[A|c < 1/2])] < 1/2$ and takes action 0. From the perspective of an outside observer, the first piece of evidence acts as a reference point and conditions how further information is interpreted. In this particular case, a high first signal eventually leads to the action optimal in state A and a low first signal eventually leads to the action optimal in B .²⁹ Naturally, the ordering of signals cannot affect beliefs or actions in a standard model where the exact signals \underline{c} and \bar{c} are processed. The result is consistent with experimental evidence according to which first impressions, acting as an anchor, matter. These observations have been made for important real life situations such as employment interview (Dougherty et al., 1994), medical diagnosis (Elstein et al., 1978) and judicial reasoning (Pennington and Hastie, 1993) among others.

In a similar vein and from the perspective of an outsider, subjects will appear *stubborn*. As they build up their confidence on one state, they become more and more likely to interpret ambiguous evidence as support for their beliefs. At the same time, strong contradictory evidence reverses their belief more dramatically. More generally, this theory supports the idea that people develop habits that are difficult to change and that people are less likely to change their mind with age.

Finally, from the perspective of an outside observer, decisions over time are interrelated. This phenomenon might be at the source of observed trends in markets, especially in financial markets, in which time series indicators are often correlated over time. As an example, it is widely documented (but poorly understood) that financial returns follow GARCH processes, where current volatility is affected by past volatility.

4.1.2. Polarization of opinions

Individuals who exhibit confirmatory biases may interpret the same information in opposite ways. The result is known as the *polarization effect*: mixed evidence is given to subjects whose existing views lie on both sides of the evidence and, as a consequence, their beliefs move farther apart. In an early work, Lord et al. (1979) presented studies on the deterrent effect of death penalty to a pool of subjects. When asked about the merits of death penalty, people who were initially in favor of (respectively against) capital punishment were more in favor (respectively against) after reading the studies (see also Darley and Gross, 1983 and Plous, 1991).

The literature explains this effect in terms of cognitive biases and non-Bayesian information processing. It is argued that individuals focus attention on the elements that support their original beliefs and (consciously or unconsciously) neglect the elements that contradict them. Our analysis suggests that attentional deficits and non-Bayesian information processing need not be at the origin of this result. Instead, the decision-threshold mechanism can fully account for this behavior.

Implication 2 (Polarization). *Individuals with different priors who receive identical information may move their beliefs farther apart.*

²⁸ As beautifully expressed by Leo Tolstoy, "the most difficult subjects can be explained to the most slow-witted man if he has not formed any idea of them already; but the simplest thing cannot be made clear to the most intelligent man if he is firmly persuaded that he knows already, without a shadow of doubt, what is laid before him." (The Kingdom of God is Within You, Chapter III).

²⁹ Rabin and Schrag (1999) obtain a similar conclusion in a behavioral non-Bayesian updating framework where it is assumed that the individual mistakenly interprets evidence against current beliefs as supporting them.

Formally, consider two individuals, i and j , with prior beliefs p_i and p_j ($< p_i$). As shown in Section 3, $x^*(p_i) < x^*(p_j)$. Suppose that ‘mixed evidence’ is released and both individuals perceive the same level of cell firing $c \in (x^*(p_i), x^*(p_j))$. Individual i will interpret the evidence in favor of A , update his belief to $p'_i > p_i$ and take a high action, whereas individual j will interpret the evidence in favor of B , update his belief to $p'_j < p_j$ and take a low action. Thus, the interpretative feature of the threshold mechanism generates a bias in the way the information is processed, which may lead to a polarized reading of identical evidence. Once again, this result cannot occur in a Bayesian world if individuals receive and process the same signal.

4.1.3. Other biases: miss error rates in visual detection tasks

When subjects are asked to visually detect targets, it is shown that miss error rates are relatively higher if the frequency of targets is low (Mackworth, 1970; Eggin and Feinstein, 1996; Wolfe et al., 2007). In other words, low probability events (a knife in an airport luggage, a tumor in a mammogram) are incorrectly evaluated relatively more often than high probability events. This effect is usually explained as a *criterion shift*, an idea consistent with our model. The decision threshold sets a criterion to interpret evidence, which is modulated by the likelihood of the event. It thus rationalizes the fact that individuals tend to see what they expect to see, and to miss what they do not expect to see.³⁰

4.2. The role of payoffs on belief updating

The results of Section 3 suggest a relationship between beliefs and payoffs: payoffs influence the decision-threshold which, in turn, affects the set of attainable ex-post beliefs. In other words, information interpretation and subsequent belief formation are shaped by the desirability of outcomes. This has a series of interesting implications.

4.2.1. Preferences and beliefs

The way information is interpreted depends on how the decision-maker feels about alternatives, which itself affects the beliefs held ex-post. To illustrate this idea in the simplest possible terms, consider a population of agents identical in all respects except for π_A , the marginal disutility of taking a low action when the state is A . This difference can be subjective (fear of dying) or objective (likelihood of recovery from an injury). Assume that $\pi_A \in (\underline{\pi}, \bar{\pi})$. All individuals have the same prior p . They receive the same information and, for simplicity, experiment the same level of neuronal cell firing c . We obtain the following result:

Implication 3 (*Preferences shape beliefs*). *Individuals with identical priors but different utilities will hold systematically different posterior beliefs.*

From Section 3, we know that $dx^*/d\pi_A < 0$. Therefore, for a given cell-firing c , there exists a cutoff $\tilde{\pi}_A$ such that $c < x^*(\pi_A)$ for all $\pi_A < \tilde{\pi}_A$ and $c > x^*(\pi_A)$ for all $\pi_A > \tilde{\pi}_A$. Individuals in the first group take action 0 and revise their belief downwards. Individuals in the second group take action 1 and revise their belief upwards. In words, subjects sharing a prior and exposed to the same evidence may end up making different choices *and* holding different opinions. So, for instance, stronger individuals who are objectively less threatened by predators will decide more often to go hunting than their weaker peers. This is rather obvious. More interestingly, they will also be more likely to believe that the environment is safe. This second conclusion cannot occur in the standard Bayesian framework where the exact c is processed. It also suggests an endogenous mechanism for what an outside observer may perceive as ‘overconfidence’ or ‘ex-post rationalization’: individuals who go hunting report low danger whereas individual who stay in the cave report high danger.

Interestingly, there is often a subjective component in payoffs. Our mechanism may account for the fact that ‘delusional’ patients tend to report high confidence in states they consider undesirable. Those beliefs are labeled as irrational in the psychology literature (Baron, 1988). According to our analysis, they may be internally consistent. For example, an individual who is terrified of dying (large perceived π_A) will set a very low decision-threshold. As a consequence, he will interpret any minor symptom as a threat to his life. A treatment will then be more effective if it is directed into convincing the patient that his subjective fear of dying is unrealistically high than if the fear of dying is left untouched and the treatment targets the (correct given his fear) interpretation of the minor symptoms.

4.2.2. Non-expected utility theory

The result emphasized above on ‘preferences shaping beliefs’ relates more generally to the idea that probabilities over outcomes and attitudes towards risks are interrelated. This issue has long been explored in the literature on non-expected utility, and several alternative utility representations have been suggested. Among the most notable, prospect theory proposes a probability weighting function to capture the tendency of decision-makers to over-react to small probability events and under-react to medium and large probability events (Kahneman and Tversky, 1979;

³⁰ We do not want to overemphasize this application since it would be also consistent with a standard Bayesian framework where attention is costly: as the likelihood of an event decreases, so does the expected benefit of looking for an error. It is empirically difficult to discriminate between these two explanations.

Prelec, 1998). Rank-dependent expected utility offers a representations in which outcomes are ranked and unlikely extreme outcomes are over-weighted (Quiggin, 1982). Finally, security-potential/aspiration theory builds on the idea that fear and hope lead individuals to overweight the probabilities attached to the most undesirable and the most desirable events, respectively (see Lopes, 1987, for the theory and Shefrin, 2008, for applications to financial decision-making).³¹ In those models, probabilities are considered objective but weighting functions introduce a subjective aspect as they may vary across individuals.³² Our model provides an alternative utility representation that accounts for similar effects.

Implication 4 (Probability functions). *An optimal decision-threshold mechanism generates payoff-dependent posterior beliefs.*

In our model, reports made on posterior beliefs are correlated with payoffs and behavior is not consistent with expected utility theory. In particular, the individual can be represented as an entity with generic utility function of the form:

$$\sum_{s=A,B} P_s(\pi_A, \pi_B, \mathcal{L}) \pi_s \hat{l}(s, \gamma),$$

where $P_s(\pi_A, \pi_B, \mathcal{L})$ is the posterior belief and $\pi_s \hat{l}(s, \gamma)$ is the loss when action γ is taken in state s . The key issue is that the desirability of outcomes affects the decision-threshold, and therefore the choices *and* posterior beliefs. Perhaps the main difference between our model and the general approach followed in decision theory is that we do not explicitly consider a distinction between objective and subjective probabilities in the brain. Instead, the individual forms an opinion based on the stimuli received by the sensory system and on the report sent to the decision system. In particular, if the paradigm is uncommon (as in the Allais paradox for example) or presented in ambiguous terms, the decision-maker may form subjective probabilities and set decision-thresholds accordingly.

4.3. Elimination strategy in complex choices

The literature in social psychology has emphasized the difficulty for individuals to think through complex decisions with many alternatives (Payne, 1982; Timmermans, 1993). To deal with these problems, individuals typically focus on a few salient options and neglect the rest. This was first discussed in the “elimination by aspects” theory of Tversky (1972), and has subsequently received strong experimental support (see e.g. Payne et al., 1988).

A three-state extension of our analysis can help understand this problem better. Suppose that $S = \{A, O, B\}$, $\Gamma = [0, 1]$ and denote by p_S the prior probability of state S . Also, $\tilde{l}(\gamma, A) = l(\gamma - 1)$, $\tilde{l}(\gamma, O) = l(\gamma - \frac{1}{2})$ and $\tilde{l}(\gamma, B) = l(\gamma - 0)$ with $l(z) = -|z|$. Finally, when $S = O$, the probability of a cell firing level c is $f(c|O)$, with $\frac{d}{dc} \frac{f(c|B)}{f(c|O)} < 0$ and $\frac{d}{dc} \frac{f(c|O)}{f(c|A)} < 0$ for all c . In words, ‘low’, ‘intermediate’ and ‘high’ cell firing is imperfect evidence of states B , O and A respectively.

Given a linear loss function, a simple extension of the argument in Lemma 1 implies that only three actions can be optimal: $\tilde{\gamma} \in \{0, \frac{1}{2}, 1\}$. An optimal process needs to set only two decision-thresholds, which we denote \underline{x} and \bar{x} , to discriminate between these three alternatives. The actions selected are 0 if $c < \underline{x}$, 1/2 if $c \in [\underline{x}, \bar{x}]$ and 1 if $c > \bar{x}$. The constrained optimal strategy of a process that can only set *one* decision-threshold is described below (and proved in Appendix A.6).

Implication 5 (Elimination). *The optimal one-threshold process discriminates perfectly between the two actions that are a priori most likely and fully disregards the third one. The efficiency loss relative to a two-thresholds process is greatest when all states are equally likely and smallest when one of the states is highly unlikely.*

A one-threshold process necessarily results in some loss, as it can only discriminate between two actions. The issue is to determine where should the threshold be set. One could think that, even if an action is left out, it will still affect how the individual discriminates between the other two. This intuition is incorrect. Instead, the process sacrifices the action which is optimal in the state most unlikely to occur (e.g., action 1 if p_A is low relative to p_O and p_B), and then discriminates *optimally* between the other two (e.g., sets threshold \underline{x} to perfectly differentiate between actions 0 and 1/2). The result is very much in line with the elimination by aspects theory discussed above. States are categorized by relevance, which is a function of how probable they are. The least relevant option is fully disregarded and the most relevant ones are evaluated optimally within the reduced set. When an option is ignored then, for the purpose of the choice to be made, it is as if it did not even exist. In other words, the heuristic of the elimination strategy is proved to be optimal in this particular case and given the processing constraints. An immediate implication is that the cost of discriminating between only two actions will be a function of how likely the third one is.

³¹ There exists a growing literature studying how the brain represents probabilities and preferences over outcomes in order to identify an appropriate model of decision-making under risk. See, for instance, Chew and Sagi (2008) for an axiomatization of the source preference hypothesis, and Hsu et al. (2005) and Chew et al. (2008) for neuroimaging studies identifying the brain regions that encode risk and uncertainty.

³² Along those lines, there is also the generalized model of certainty equivalence developed by Chew (1983) which satisfies some desirable properties of expected utility (transitivity, betweenness, stochastic dominance) and, at the same time, resolves the Allais paradox.

5. Conclusion

Building theoretical models of brain processes is an important step both for economics and neuroscience. For economics, incorporating *physiological costs and constraints* in the capacity of individuals to evaluate situations, process information and reach conclusions has two advantages. First, it provides guidelines regarding the plausibility of different assumptions when we try and model bounded rationality. Second, it provides micro-microfoundations for some well-documented biases in choices.³³ For neuroscience, formal models of the brain can provide testable implications about the functionality of different brain systems and inspire new experiments (see Brocas and Carrillo, 2008, for a more detailed discussion of this methodology). This paper has taken one step in that direction. We have provided a theoretical framework to study information processing in the brain. We have then used this framework to predict decisions in behaviorally relevant environments. Last, we have analyzed several implications of our theory and shown that our predictions are consistent with some anomalies and biases documented in psychology. Interestingly, those biases all originate in the same physiological constraints, and are thus likely to be observed in conjunction.

Our model can be extended in a number of directions. First and unlike the diffusion model, we do not account for the temporal aspects of information processing. In our theory, it is simply assumed that the threshold is reached. Correlating reaction times, experimental conditions and likelihood of mistakes in an optimization model would be of significant value. Second, we could analyze situations where individuals learn also from their choices. In particular, we could compare the difference in beliefs before and after a decision is made and a stochastic outcome is realized. This would shed light on the reaction to expected and unexpected events. Last, a natural extension would be to analyze how individuals take decisions in strategic settings. For instance, we could measure the impact of confirmatory biases on decisions when agents engage in synergistic or competitive activities.

Appendix A

A.1. Proof of Proposition 1

Assume that the parameters of the model are such that $\frac{f(0|B)}{f(0|A)} > \frac{p}{1-p} \frac{\pi_A}{\pi_B} > \frac{f(1|B)}{f(1|A)}$.³⁴ Taking the first-order condition in (5), we find that $x^*(p)$ satisfies (6). Given **(A1)**, $x^*(p)$ is unique and the local second-order condition is satisfied:

$$\begin{aligned} \frac{\partial^2 V}{\partial x^2} \Big|_{x^*} &= -pf'(x^*|A)\pi_A(l(0) - l(1)) + (1-p)f'(x^*|B)\pi_B(l(0) - l(1)) \\ &= (1-p)\pi_B f(x^*|A)(l(0) - l(1)) \left(\frac{f(x^*|B)}{f(x^*|A)} \right)' < 0. \end{aligned}$$

$dx^*/dp < 0$ is immediate from **(A1)**. Finally, we can ex-post check that it is optimal to select $\gamma^* = 1$ when $c > x^*$ and $\gamma^* = 0$ when $c < x^*$. We have:

$$\begin{aligned} \Pr(A|c > x^*) &> \frac{\pi_B}{\pi_A + \pi_B} \Leftrightarrow p\pi_A(1 - F(x^*|A)) > (1-p)\pi_B(1 - F(x^*|B)) \Leftrightarrow \frac{1 - F(x^*|A)}{f(x^*|A)} > \frac{1 - F(x^*|B)}{f(x^*|B)} \\ \Pr(A|c < x^*) &< \frac{\pi_B}{\pi_A + \pi_B} \Leftrightarrow p\pi_A F(x^*|A) > (1-p)\pi_B F(x^*|B) \Leftrightarrow \frac{F(x^*|A)}{f(x^*|A)} < \frac{F(x^*|B)}{f(x^*|B)} \end{aligned}$$

Both inequalities are satisfied given **(A1)**. This completes the proof.

A.2. Proof of Proposition 2

Given (9), we can rewrite the value function as:

$$\begin{aligned} V(x) &= p(1 - F(x|A))\pi_A l(1 - \gamma^{**}(\bar{p}(x))) + (1-p)(1 - F(x|B))\pi_B l(\gamma^{**}(\bar{p}(x))) \\ &\quad + pF(x|A)\pi_A l(1 - \gamma^{**}(\underline{p}(x))) + (1-p)F(x|B)\pi_B l(\gamma^{**}(\underline{p}(x))) \end{aligned} \quad (15)$$

The optimal threshold maximizes (15) given (7) and (8). The first-order condition is:

$$\frac{\partial V(x)}{\partial x} \Big|_{x=x^{**}} = 0 \Rightarrow \frac{f(x^{**}|B)}{f(x^{**}|A)} = \frac{p}{1-p} \frac{\pi_A}{\pi_B} \frac{l(1 - \gamma^{**}(\bar{p}(x^{**}))) - l(1 - \gamma^{**}(\underline{p}(x^{**})))}{l(\gamma^{**}(\underline{p}(x^{**}))) - l(\gamma^{**}(\bar{p}(x^{**})))} \quad (16)$$

³³ One could draw a parallel with the theory of organizations, where a more accurate modeling of firm constraints (agency problems, restricted information channels, limited resource allocation) has helped understanding organizational choices.

³⁴ This condition ensures that the optimal solution is interior. If it is not satisfied, then the threshold is set at a bound. This means that the threshold is either always surpassed ($x^* = 0$) or never surpassed ($x^* = 1$). In either case, no new information is acquired and the prior coincides with the posterior. This uninteresting case corresponds to the situation where the prior beliefs are so skewed towards one alternative that no signal could ever reverse them.

From (16) and using (7) and (8), we get:

$$\begin{aligned} \left. \frac{\partial^2 V(x; p)}{\partial x \partial p} \right|_{x=x^{**}} &= -\frac{\pi_B}{p} \left[f(x^{**}|B)(l(\gamma^{**}(\underline{p})) - l(\gamma^{**}(\bar{p}))) + F(x^{**}|B)l'(\gamma^{**}(\underline{p})) \frac{d\gamma^{**}(\underline{p})}{dx} \right]_{x^{**}} \\ &\quad + (1 - F(x^{**}|B))l'(\gamma^{**}(\bar{p})) \frac{d\gamma^{**}(\bar{p})}{dx} \Big|_{x^{**}} \Big] \\ &= -\frac{\pi_B}{p} \times \frac{d}{dx} [F(x|B)l(\gamma^{**}(\underline{p})) + (1 - F(x|B))l(\gamma^{**}(\bar{p}))]_{x=x^{**}} \end{aligned}$$

Similarly,

$$\begin{aligned} \left. \frac{\partial^2 V(x; p)}{\partial x^2} \right|_{x=x^{**}} &= (1 - p)\pi_B \frac{\left(\frac{f(x|B)}{f(x|A)}\right)'}{\frac{f(x|B)}{f(x|A)}} \left[f(x|B)(l(\gamma^{**}(\underline{p})) - l(\gamma^{**}(\bar{p}))) \right. \\ &\quad + F(x|B)l'(\gamma^{**}(\underline{p})) \frac{d\gamma^{**}(\underline{p})}{dx} \frac{\left(\frac{f(x|B)}{f(x|A)}\right)'}{\left(\frac{f(x|B)}{f(x|A)}\right)'} \frac{\left(\frac{F(x|B)}{F(x|A)}\right)'}{\frac{F(x|B)}{F(x|A)}} \\ &\quad \left. + (1 - F(x|B))l'(\gamma^{**}(\bar{p})) \frac{d\gamma^{**}(\bar{p})}{dx} \frac{\left(\frac{f(x|B)}{f(x|A)}\right)'}{\left(\frac{f(x|B)}{f(x|A)}\right)'} \frac{\left(\frac{1-F(x|B)}{1-F(x|A)}\right)'}{\frac{1-F(x|B)}{1-F(x|A)}} \right]_{x=x^{**}} \end{aligned}$$

By (A2), $\frac{\left(\frac{f(x|B)}{f(x|A)}\right)'}{\left(\frac{f(x|B)}{f(x|A)}\right)'} \frac{\left(\frac{F(x|B)}{F(x|A)}\right)'}{\frac{F(x|B)}{F(x|A)}} \leq 1$ and $\frac{\left(\frac{f(x|B)}{f(x|A)}\right)'}{\left(\frac{f(x|B)}{f(x|A)}\right)'} \frac{\left(\frac{1-F(x|B)}{1-F(x|A)}\right)'}{\frac{1-F(x|B)}{1-F(x|A)}} \leq 1$. Therefore, $\frac{\partial^2 V(x; p)}{\partial x \partial p} \Big|_{x=x^{**}} < 0 \Rightarrow \frac{\partial^2 V(x; p)}{\partial x^2} \Big|_{x=x^{**}} < 0$ and the proposition follows.

A.2.1. Characterization of the equilibrium with quadratic loss

Suppose that $l(z) = \alpha - \beta z^2$ with $\beta > 0$ and $\pi_A = \pi_B = 1$. Under this restriction, (7) and (8) become:

$$\gamma^{**}(\bar{p}(x)) = \bar{p}(x) \quad \text{and} \quad \gamma^{**}(\underline{p}(x)) = \underline{p}(x)$$

Therefore, (10) has the following simple expression:

$$\frac{f(x^{**}|B)}{f(x^{**}|A)} = \frac{p}{1-p} \frac{(1 - \bar{p}(x^{**})) + (1 - \underline{p}(x^{**}))}{\bar{p}(x^{**}) + \underline{p}(x^{**})} \tag{17}$$

Let $P \equiv \frac{p}{1-p}$. The F.O.C. (17) can be rewritten as:

$$\left. \frac{1-p}{p^2} \frac{\partial V(x; p)}{\partial x} \right|_{x^{**}} = k(x^{**}, P)$$

where

$$\begin{aligned} k(x, P) &\equiv \left(\frac{1 - F(x|A)}{(1 - F(x|A))P + (1 - F(x|B))} - \frac{F(x|A)}{F(x|A)P + F(x|B)} \right) \\ &\quad \times \left[f(x|B) \left(\frac{1 - F(x|A)}{(1 - F(x|A))P + (1 - F(x|B))} + \frac{F(x|A)}{F(x|A)P + F(x|B)} \right) \right. \\ &\quad \left. - f(x|A) \left(\frac{1 - F(x|B)}{(1 - F(x|A))P + (1 - F(x|B))} + \frac{F(x|B)}{F(x|A)P + F(x|B)} \right) \right]. \end{aligned}$$

Differentiating the F.O.C., we get

$$\left. \frac{\partial^2 V(x; p)}{\partial x \partial p} \right|_{x^{**}} \propto \frac{\partial k(x^{**}, P)}{\partial P}$$

After some tedious algebra, and using (17), we get:

$$\frac{\partial k(x^{**}, P)}{\partial P} = -f(x^{**}|A) \frac{\left(\frac{F(x^{**}|B)}{F(x^{**}|A)} - \frac{1-F(x^{**}|B)}{1-F(x^{**}|A)}\right)^3}{\left(2P + \frac{F(x^{**}|B)}{F(x^{**}|A)} + \frac{1-F(x^{**}|B)}{1-F(x^{**}|A)}\right) \left(P + \frac{F(x^{**}|B)}{F(x^{**}|A)}\right)^2 \left(P + \frac{1-F(x^{**}|B)}{1-F(x^{**}|A)}\right)^2} < 0$$

which guarantees that $dx^{**}/dp < 0$ is satisfied in every locally optimal threshold.

A.2.2. Analytical example in the linear and quadratic cases

Let $F(c|A) = c^2$ and $F(c|B) = c$. From (6) and (17) and after some algebra, the optimal thresholds with linear ($l(z) = -|z|$) and quadratic ($l(z) = -z^2$) payoffs are respectively:

$$x^*(p) = \frac{1-p}{2p} \quad \text{and} \quad x^{**}(p) = \frac{\sqrt{1-p}}{\sqrt{1-p} + \sqrt{1+p}}$$

where x^* and x^{**} are interior if $p > 1/3$. In this example, the optimal threshold is always less extreme with quadratic than with linear payoffs: $x^* \geq x^{**} \geq 1/3$ for all $p \leq 3/5$.

A.3. Proof of Proposition 3

Taking the first-order condition in (13) and applying the envelope theorem, we get:

$$\frac{\partial W(y; p)}{\partial y} \Big|_{y=y^*} = 0 \quad \Rightarrow \quad \frac{f_1(y^*|B)}{f_1(y^*|A)} = \frac{p}{1-p} \frac{F_2(x^*(\underline{p}(y^*))|A) - F_2(x^*(\bar{p}(y^*))|A)}{F_2(x^*(\underline{p}(y^*))|B) - F_2(x^*(\bar{p}(y^*))|B)} \frac{\pi_A}{\pi_B}$$

The rest of the proof follows the exact same steps as the proof of Proposition 2 and is therefore omitted for the sake of brevity.

A.3.1. Equilibrium with second stage linear densities

Let $g_2(c) = 2c$ and $f_2(c) = 2(1-c)$, $\pi_A = \pi_B = 1$, and let us keep a general formulation for the first stage cell firing densities. After some algebra, the first-order condition (14) can be rewritten as:

$$\frac{f_1(y^*|B)}{f_1(y^*|A)} = \frac{p}{1-p} \frac{(1 - \bar{p}(y^*)) + (1 - \underline{p}(y^*))}{\bar{p}(y^*) + \underline{p}(y^*)}$$

which is exactly the same expression as (17), and the result follows.

Example 1 (Continued). Suppose $l(0) = 1$, $l(1) = 0$ and $p = 1/2$. It follows that $y^* = 1/2$, $\bar{p}(1/2) = 3/4$, $x^*(\bar{p}(1/2)) = 1/4$, $\underline{p}(1/2) = 1/4$, and $x^*(\underline{p}(1/2)) = 3/4$. A continuum of thresholds in stage 1 is formally equivalent to observing the exact cell firing c_1 . Given $p(A|c_1) = c_1$, the expected loss under that process is:

$$\tilde{W} = p \int_0^1 \Pr(c_1|A) [1 - F(x^*(p(A|c_1))|A)] dc_1 + (1-p) \int_0^1 \Pr(c_1|B) [F(x^*(p(A|c_1))|B)] dc_1 = \frac{5}{6}$$

Suppose instead that the brain uses a process with only one threshold. The individual at stage 1 only learns whether $c_1 \geq y^*$ ($= 1/2$). Following (13), his expected loss is:

$$W = \Pr(A) \left[\Pr\left(c_1 > \frac{1}{2} \Big| A\right) \Pr\left(c_1 > \frac{1}{4} \Big| A\right) + \Pr\left(c_1 < \frac{1}{2} \Big| A\right) \Pr\left(c_1 > \frac{3}{4} \Big| A\right) \right] \\ + \Pr(B) \left[\Pr\left(c_1 < \frac{1}{2} \Big| B\right) \Pr\left(c_1 < \frac{3}{4} \Big| B\right) + \Pr\left(c_1 > \frac{1}{2} \Big| B\right) \Pr\left(c_1 < \frac{1}{4} \Big| B\right) \right] = \frac{13}{16} < \tilde{W}$$

In this example, the expected utility loss of using the one-threshold process is only 2.5%.

A.4. Optimal threshold with two actions and a continuum of states

Let $S = [0, 1]$ and $\Gamma = \{0, 1\}$. In our example, $s \in S$ captures the proportion of predators in the neighborhood. The individual decides between hunting ($\gamma = 0$) and staying in the cave ($\gamma = 1$). We order the states by the increasing degree of danger, from safest ($s = 0$) to most dangerous ($s = 1$). Payoffs are $\tilde{l}(1, s) = \pi_s l(1-s)$ and $\tilde{l}(0, s) = \pi_s l(-s)$, where $l(z) = l(-z)$ for all z and $l'(z) < 0$ for all $z > 0$. The probability of cell firing level c given state s is now $f(c|s)$. The generalization of MLRP to the continuous case is:

Assumption 1' (Continuous MLRP).

$$\frac{d}{dc} \left(\frac{f_s(c|s)}{f(c|s)} \right) \geq 0 \quad \text{for all } c \text{ and } s. \tag{A1'}$$

The individual believes that the state is s with probability $p(s)$. The expected payoff (3) is generalized as $L(p(s), \gamma) = \int_0^1 p(s)\pi_s l(\gamma - s) ds$ and the optimal action is:

$$\hat{\gamma} = 1 \quad \text{if} \quad \int_0^1 p(s)\pi_s(l(1-s) - l(s)) ds > 0 \quad \text{and} \quad \hat{\gamma} = 0 \quad \text{if} \quad \int_0^1 p(s)\pi_s(l(1-s) - l(s)) ds < 0$$

To simplify the analysis, assume that $\pi_s = \pi_l$ for all states $s < 1/2$ and $\pi_s = \pi_u$ for all states $s \geq 1/2$. The value function is:

$$\begin{aligned} V(x; p(s)) &= \Pr(c > x)L(p(s|c > x), 1) + \Pr(c < x)L(p(s|c < x), 0) \\ &= \int_0^1 p(s)\pi_s((1 - F(c|s))l(1-s) + F(c|s)l(s)) ds \end{aligned} \tag{18}$$

Denote by $\hat{x}(p(s)) = \arg \max_x V(x; p(s))$. Taking the F.O.C. in (18), we obtain:

$$-\int_0^1 p(s)\pi_s f(\hat{x}|s)(l(1-s) - l(s)) ds = 0 \tag{19}$$

The local S.O.C. is:

$$\begin{aligned} \frac{\partial^2 V}{\partial x^2} \Big|_{\hat{x}} &= -\int_0^1 p(s)\pi_s f_x(\hat{x}|s)(l(1-s) - l(s)) ds \\ &= \int_0^1 \left(-\frac{f_x(\hat{x}|s)}{f(\hat{x}|s)} \right) p(s)\pi_s f(\hat{x}|s) (l(1-s) - l(s)) ds \end{aligned}$$

Let $h(s) \equiv -\frac{f_x(\hat{x}|s)}{f(\hat{x}|s)}$. By **(A1')**, $h'(s) \leq 0$. We can then rewrite the local S.O.C. as:

$$\begin{aligned} \frac{\partial^2 V}{\partial x^2} \Big|_{\hat{x}} &= \int_0^{1/2} h(s)p(s)\pi_s f(\hat{x}|s)(l(1-s) - l(s)) ds + \int_{1/2}^1 h(s)p(s)\pi_s f(\hat{x}|s)(l(1-s) - l(s)) ds \\ &< h(1/2) \left[\int_0^{1/2} p(s)\pi_s f(\hat{x}|s)(l(1-s) - l(s)) ds + \int_{1/2}^1 p(s)\pi_s f(\hat{x}|s)(l(1-s) - l(s)) ds \right] = 0 \end{aligned}$$

which means that the threshold $\hat{x}(p(s))$ defined by (19) is indeed a unique maximum.

Suppose now that $(\frac{p(s)}{q(s)})' \leq 0$, then:

$$\begin{aligned} \frac{\partial V(x; p(s))}{\partial x} \Big|_{\hat{x}(p(s))} &= -\int_0^1 \left(\frac{p(s)}{q(s)} \right) q(s)\pi_s f(\hat{x}(p(s))|s)(l(1-s) - l(s)) ds \\ &> -\left(\frac{p(1/2)}{q(1/2)} \right) \int_0^1 q(s)\pi_s f(\hat{x}(p(s))|s)(l(1-s) - l(s)) ds \end{aligned}$$

Therefore,

$$\frac{\partial V(x; p(s))}{\partial x} \Big|_{\hat{x}(p(s))} = 0 > \left(\frac{p(1/2)}{q(1/2)} \right) \frac{\partial V(x; q(s))}{\partial x} \Big|_{\hat{x}(p(s))} \Rightarrow \hat{x}(p(s)) > \hat{x}(q(s))$$

In words, if one individual puts more weight in higher states than another one in a MLRP sense, $(\frac{q(s)}{p(s)})' \geq 0$, then he also sets a lower threshold. This property is simply a generalization of the comparative statics on p to the case of a continuous distribution of beliefs.

Finally, we need to check that it is indeed optimal to choose $\hat{y} = 1$ when $c > \hat{x}$ and $\hat{y} = 0$ when $c < \hat{x}$. Let $\mathcal{J}(x) \equiv L(p(s|c = x), 1) - L(p(s|c = x), 0)$, also $p(s|c = x) \equiv j(s|x) = \frac{p(s)f_x(x|s)}{\int_0^1 p(s)f_x(x|s) ds}$ and $J(s|x) = \int_0^s j(\tilde{s}|x) d\tilde{s}$. We use the fact that $\pi_s = \pi_l$ for all $s < 1/2$ and $\pi_s = \pi_u$ for all $s \geq 1/2$. Integrating by parts:

$$\begin{aligned} \mathcal{J}(x) &= \int_0^1 j(s|x)\pi_s(l(1-s) - l(s)) ds \\ &= \pi_u(l(0) - l(1)) + \int_0^1 J(s|x)\pi_s(l'(1-s) + l'(s)) ds \end{aligned}$$

Therefore

$$\frac{d\mathcal{J}(x)}{dx} = \int_0^1 J_x(s|x)\pi_s(l'(1-s) + l'(s)) ds > 0$$

since, by **(A1')**, we know that $F_s(x|s) < 0$ and therefore $J_x(s|x) < 0$. From (19), $\mathcal{J}(\hat{x}) = 0$, so $\mathcal{J}(x) \geq 0$ for all $x \geq \hat{x}$. This also proves that, for the purpose of the action to be taken, it is equivalent to learn c or to learn whether c is greater or smaller than \hat{x} .

Last, setting $\hat{\mathcal{L}} = \frac{\pi_u}{\pi_l}$ and differentiating (19) with respect to $\hat{\mathcal{L}}$, we obtain:

$$\frac{1}{\pi_l} \frac{\partial^2 V}{\partial x^2} \Big|_{\hat{x}} \frac{\partial \hat{x}}{\partial \hat{\mathcal{L}}} - \int_{1/2}^1 p(s) f_x(\hat{x}|s) (l(1-s) - l(s)) ds = 0$$

which implies that $d\hat{x}/d\hat{\mathcal{L}} < 0$. Summing up, the conclusions stated in Proposition 1 (no loss of utility by setting only one optimal cutoff and comparative statics of optimal cutoff with respect to prior beliefs and cost of wrong actions) extend to the case of two actions and a continuum of states.

A.5. Optimal thresholds in simple vs. complex activities

Consider activities “ α ” and “ β ” such that low cell firing when $S = A$ or high cell firing when $S = B$ are uniformly less frequent in α -activities than in β -activities. Formally:

$$\frac{\partial}{\partial c} \left(\frac{f_\alpha(c|A)}{f_\beta(c|A)} \right) > 0 \quad \text{and} \quad \frac{\partial}{\partial c} \left(\frac{f_\alpha(c|B)}{f_\beta(c|B)} \right) < 0 \tag{20}$$

where $f_k(c|S)$ is the probability of cell firing c in situation $k \in \{\alpha, \beta\}$ given state S . The idea is simply that “neuronal mistakes”, defined as low cell firing when $S = A$ or high cell firing when $S = B$, are uniformly less frequent in α -activities than in β -activities. In other words, α activities represent simple (concrete, common, temporally close) choices whereas β activities represent complex (abstract, exceptional, temporally distant) choices. For technical reasons, we also assume that **(A1)** is satisfied in both type of activities.

$\left(\frac{f_\alpha(c|A)}{f_\beta(c|A)}\right)' > 0$, $\left(\frac{f_\beta(c|B)}{f_\beta(c|A)}\right)' < 0$, $\left(\frac{f_\alpha(c|B)}{f_\beta(c|B)}\right)' < 0$, $\left(\frac{f_\alpha(c|A)}{f_\alpha(c|B)}\right)' < 0 \Rightarrow \frac{f'_\alpha(c|A)}{f_\alpha(c|A)} > \frac{f'_\beta(c|A)}{f_\beta(c|A)} > \frac{f'_\beta(c|B)}{f_\beta(c|B)} > \frac{f'_\alpha(c|B)}{f_\alpha(c|B)}$. Now, suppose there exists $\hat{c} \in (0, 1)$ such that $\frac{f_\alpha(\hat{c}|B)}{f_\alpha(\hat{c}|A)} = \frac{f_\beta(\hat{c}|B)}{f_\beta(\hat{c}|A)}$. Then,

$$\frac{d}{dc} \left[\frac{f_\alpha(c|B)}{f_\alpha(c|A)} - \frac{f_\beta(c|B)}{f_\beta(c|A)} \right]_{c=\hat{c}} = \frac{f_\alpha(\hat{c}|B)}{f_\alpha(\hat{c}|A)} \left(\frac{f'_\alpha(\hat{c}|B)}{f_\alpha(\hat{c}|B)} - \frac{f'_\alpha(\hat{c}|A)}{f_\alpha(\hat{c}|A)} \right) - \frac{f_\beta(\hat{c}|B)}{f_\beta(\hat{c}|A)} \left(\frac{f'_\beta(\hat{c}|B)}{f_\beta(\hat{c}|B)} - \frac{f'_\beta(\hat{c}|A)}{f_\beta(\hat{c}|A)} \right) < 0$$

so $\frac{f_\alpha(c|B)}{f_\alpha(c|A)}$ and $\frac{f_\beta(c|B)}{f_\beta(c|A)}$ cross at most once. Also, $\left(\frac{f_\alpha(c|A)}{f_\beta(c|A)}\right)' > 0$ and $\left(\frac{f_\alpha(c|B)}{f_\beta(c|B)}\right)' < 0 \Rightarrow \frac{f_\alpha(0|B)}{f_\alpha(0|A)} > \frac{f_\beta(0|B)}{f_\beta(0|A)}$ and $\frac{f_\alpha(1|B)}{f_\alpha(1|A)} < \frac{f_\beta(1|B)}{f_\beta(1|A)}$.

Together with the previous result, it means that there exists $\hat{x} \in (0, 1)$ such that $\frac{f_\alpha(x|B)}{f_\alpha(x|A)} \geq \frac{f_\beta(x|B)}{f_\beta(x|A)}$ for all $x \leq \hat{x}$. Denote by

$x_k^*(p)$ the optimal threshold in activity k as a function of p . Given (6), there exists \hat{p} such that $\frac{f_\alpha(x_\alpha^*(\hat{p})|B)}{f_\alpha(x_\alpha^*(\hat{p})|A)} = \frac{f_\beta(x_\beta^*(\hat{p})|B)}{f_\beta(x_\beta^*(\hat{p})|A)} = \frac{\hat{p}}{1-\hat{p}}$,

that is, $x_\alpha^*(\hat{p}) = x_\beta^*(\hat{p}) = x^*(\hat{p}) \equiv \hat{x}$. For all $p \geq \hat{p}$, $\frac{f_\alpha(x_\alpha^*(p)|B)}{f_\alpha(x_\alpha^*(p)|A)} = \frac{f_\beta(x_\beta^*(p)|B)}{f_\beta(x_\beta^*(p)|A)} = \frac{p}{1-p} \Rightarrow x_\beta^*(p) \leq x_\alpha^*(p) \leq \hat{x}$. Overall, optimal thresholds are more sensitive to initial beliefs in complex than in simple activities.

Now, suppose $p < \hat{p}$. Then, $x_\alpha^* < x_\beta^*$. Given **(A1)**, $\frac{F_\alpha(x_\alpha^*|A)}{F_\alpha(x_\alpha^*|B)} < \frac{F_\alpha(x_\beta^*|A)}{F_\alpha(x_\beta^*|B)}$. Given (20), $\frac{F_\alpha(x_\beta^*|A)}{F_\alpha(x_\beta^*|B)} < \frac{F_\beta(x_\beta^*|A)}{F_\beta(x_\beta^*|B)}$. Finally, $\frac{F_\alpha(x_\alpha^*|A)}{F_\alpha(x_\alpha^*|B)} <$

$$\frac{F_\beta(x_\beta^*|A)}{F_\beta(x_\beta^*|B)} \Leftrightarrow \Pr_\beta[A|0] > \Pr_\alpha[A|0].$$

Analogously, if $p > \hat{p}$ then $x_\alpha^* > x_\beta^*$. Given (A1) and (20), $\frac{1-F_\beta(x_\beta^*|B)}{1-F_\beta(x_\beta^*|A)} > \frac{1-F_\beta(x_\alpha^*|B)}{1-F_\beta(x_\alpha^*|A)} > \frac{1-F_\alpha(x_\alpha^*|B)}{1-F_\alpha(x_\alpha^*|A)}$, and therefore $\Pr_\beta[B|1] > \Pr_\alpha[B|1]$. This proves that the individual is more likely to make mistakes in complex rather than simple activities.

A.6. Proof of Implication 5

A.6.1. Two-thresholds process

A necessary condition for cutoffs x_1 and x_2 ($> x_1$) to be optimal is $\tilde{\gamma} = 0$ if $c \in [0, x_1)$, $\tilde{\gamma} = \frac{1}{2}$ if $c \in [x_1, x_2]$ and $\tilde{\gamma} = 1$ if $c \in (x_2, 1]$. The value function is then:

$$\begin{aligned} V(x_1, x_2) &= \Pr(c < x_1)L(p(\cdot|c < x_1), 0) + \Pr(c \in [x_1, x_2])L\left(\frac{1}{2}; p(\cdot|c \in [x_1, x_2])\right) + \Pr(c > x_2)L(p(\cdot|c > x_2), 1) \\ &= -p_B \left[(1 - F(x_2|B)) + \frac{1}{2}(F(x_2|B) - F(x_1|B)) \right] - p_O \left[\frac{1}{2}(1 - F(x_2|O)) + \frac{1}{2}F(x_1|O) \right] \\ &\quad - p_A \left[\frac{1}{2}(F(x_2|A) - F(x_1|A)) + F(x_1|A) \right] \end{aligned} \tag{21}$$

Taking F.O.C. in (21), we obtain \underline{x} and \bar{x} . They solve:

$$\frac{f(\underline{x}|B)}{f(\underline{x}|A)} = \frac{p_A}{p_B} + \frac{p_O}{p_B} \frac{f(\underline{x}|O)}{f(\underline{x}|A)} \quad \text{and} \quad \frac{f(\bar{x}|B)}{f(\bar{x}|A)} = \frac{p_A}{p_B} - \frac{p_O}{p_B} \frac{f(\bar{x}|O)}{f(\bar{x}|A)}$$

Notice that $\underline{x} < \bar{x}$ for all $p_A, p_O, p_B \in (0, 1)^3$ and $\underline{x} = x^* = \bar{x}$ when $p_O = 0$. We also have $\frac{\partial^2 V(x_1, x_2)}{\partial x_1^2} |_{\underline{x}} = -\frac{1}{2} p_O f(\underline{x}|B) \left(\frac{f(\underline{x}|O)}{f(\underline{x}|B)} \right)' - \frac{1}{2} p_A f(\underline{x}|B) \left(\frac{f(\underline{x}|A)}{f(\underline{x}|B)} \right)' < 0$, $\frac{\partial^2 V(x_1, x_2)}{\partial x_2^2} |_{\bar{x}} = \frac{1}{2} p_O f(\bar{x}|A) \left(\frac{f(\bar{x}|O)}{f(\bar{x}|A)} \right)' + \frac{1}{2} p_B f(\bar{x}|A) \left(\frac{f(\bar{x}|B)}{f(\bar{x}|A)} \right)' < 0$, and $\frac{\partial^2 V(x_1, x_2)}{\partial x_1 \partial x_2} = 0$. Therefore \underline{x} and \bar{x} are maxima. Last, it can be easily checked that $\Pr(S|c \in \mathcal{Y})$, $S \in \{A, O, B\}$ are such that $\tilde{\gamma} = 0$ if $\mathcal{Y} = [0, \underline{x}]$, $\tilde{\gamma} = \frac{1}{2}$ if $\mathcal{Y} = [\underline{x}, \bar{x}]$, and $\tilde{\gamma} = 1$ if $\mathcal{Y} = (\bar{x}, 1]$ are indeed optimal.

A.6.2. One-threshold process

Let \check{x} be the cutoff that solves:

$$\frac{f(\check{x}|B)}{f(\check{x}|A)} = \frac{p_A}{p_B}$$

It is immediate to see that $\check{x} \in (\underline{x}, \bar{x})$. The three candidates for optimal cutoffs are:

$$\begin{cases} x_a & \text{so that } \tilde{\gamma} = 0 \text{ if } c < x_a \text{ and } \tilde{\gamma} = 1 \text{ if } c > x_a \\ x_b & \text{so that } \tilde{\gamma} = 0 \text{ if } c < x_b \text{ and } \tilde{\gamma} = 1/2 \text{ if } c > x_b \\ x_c & \text{so that } \tilde{\gamma} = 1/2 \text{ if } c < x_c \text{ and } \tilde{\gamma} = 1 \text{ if } c > x_c \end{cases}$$

These cutoffs are formally defined by:

$$\begin{cases} x_a = \arg \max_x V^a(x) \equiv \Pr(c < x)L(0; p(\cdot|c < x)) + \Pr(c > x)L(1; p(\cdot|c < x)) \\ x_b = \arg \max_x V^b(x) \equiv \Pr(c < x)L(0; p(\cdot|c < x)) + \Pr(c > x)L\left(\frac{1}{2}; p(\cdot|c < x)\right) \\ x_c = \arg \max_x V^c(x) \equiv \Pr(c < x)L\left(\frac{1}{2}; p(\cdot|c < x)\right) + \Pr(c > x)L(1; p(\cdot|c < x)) \end{cases}$$

It is straightforward to check that $x_a = \check{x}$, $x_b = \underline{x}$, $x_c = \bar{x}$. Now, fix p_O . Differentiating each first-order condition with respect to p_B , we get:

$$\frac{dx_a}{dp_B} > 0, \quad \frac{dx_b}{dp_B} > 0, \quad \frac{dx_c}{dp_B} > 0$$

Furthermore:

$$\begin{aligned} \frac{dV^a(x^a)}{dp_B} &= F(x^a|B) + F(x^a|A) - 1 \geq 0, & \frac{d^2 V^a(x^a)}{dp_B^2} &= [f(x^a|B) + f(x^a|A)] \frac{dx_a}{dp_B} \geq 0 \\ \frac{dV^b(x^b)}{dp_B} &= \frac{F(x^b|B) + F(x^b|A)}{2} \geq 0, & \frac{dV^c(x^c)}{dp_B} &= \frac{F(x^c|B) + F(x^c|A)}{2} - 1 \leq 0 \end{aligned}$$

Also, $\lim_{p_B \rightarrow 0} V^a(x^a) = -\frac{p_O}{2} < \lim_{p_B \rightarrow 0} V^c(x^c)$ and $\lim_{p_B \rightarrow 1-p_O} V^a(x^a) = -\frac{p_O}{2} < \lim_{p_B \rightarrow 1-p_O} V^b(x^b)$.

Combining these results, we have that there exist p^* such that x^c dominates x^b if $p_B < p^*$ and x^b dominates x^c if $p_B > p^*$. Also, there exist p^{**} and p^{***} such that x^c dominates x^a if $p_B < p^{**}$ and x^b dominates x^a if $p_B > p^{***}$. The ranking between p^* , p^{**} and p^{***} depend on the relative values of p_O and p_A .

References

- Abdi, H., 1994. A neural network primer. *J. Biol. Systems* 2, 247–281.
- Barlow, H., 2001. Redundancy reduction revisited. *Network: Computat. Neural Systems* 12, 241–253.
- Baron, J., 1988. *Thinking and Deciding*. Cambridge University Press, New York.
- Bogacz, R., 2007. Optimal decision-making theories: Linking neurobiology with behaviour. *Trends Cognitive Sci.* 11, 118–125.
- Brocas, I., 2009. Information processing and decision-making: evidence from the brain sciences and implications for economics. Mimeo (USC).
- Brocas, I., Carrillo, J.D., 2008. Theories of the mind. *Amer. Econ. Rev. Pap. Proc.* 98 (2), 175–180.
- Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C., Wilson, M.A., 1998. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J. Neurosci.* 18, 7411–7425.
- Calvert, R.L., 1985. The value of biased information: A rational choice model of political advice. *J. Politics* 47, 530–555.
- Camerer, C., Loewenstein, G., Prelec, D., 2005. Neuroeconomics: How neuroscience can inform economics. *J. Econ. Lit.* 43, 9–64.
- Chew, S.H., 1983. A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica* 51 (4), 1065–1092.
- Chew, S.H., Sagi, J.S., 2008. Small worlds: Modeling attitudes toward sources of uncertainty. *J. Econ. Theory* 139 (1), 1–24.
- Chew, S.H., Li, K.K., Chark, R., Songfa, Z., 2008. Source preference and ambiguity aversion: Models and evidence from behavioral and neuroimaging experiments. In: *Advances in Health Economics and Health Services Research*, vol. 20, pp. 179–201.
- Darley, J., Gross, P., 1983. A Hypothesis-confirming bias in labeling effects. *J. Person. Social Psych.* XLIV, 20–33.
- Dayan, P., Abbott, L.F., 2005. *Theoretical Neuroscience Computational and Mathematical Modeling of Neural Systems*. MIT Press, Cambridge, MA.
- Deneve, S., Latham, P., Pouget, A., 1999. Reading population codes: A neural implementation of ideal observers. *Nature Neurosci.* 2 (8), 740–745.
- Ditterich, J., Mazurek, M., Shadlen, M., 2003. Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neurosci.* 6 (8), 891–898.
- Dixit, A., Pindyck, R., 1994. *Investment Under Uncertainty*. Princeton University Press, Princeton.
- Dougherty, T.W., Turban, D.B., Calender, J.C., 1994. Confirming first impressions in the employment interview: A field study of interviewer behavior. *J. Appl. Psych.* 79, 659–665.
- Eggin, T.K., Feinstein, A.R., 1996. Context bias: A problem in diagnostic radiology. *J. Amer. Med. Assoc.* 276, 1752–1755.
- Elstein, A.S., Shulman, L.S., Sprafka, S.A., 1978. *Medical Problem Solving: An Analysis of Clinical Reasoning*. Harvard University Press, Cambridge, MA.
- Ernst, M.O., Banks, M., 2002. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–433.
- Fehr, E., Camerer, C., 2007. Social neuroeconomics: The neural circuitry of social preferences. *Trends Cognitive Sci.* 11 (10), 419–426.
- Fryer, R., Jackson, M.O., 2008. A categorical model of cognition and biased decision making. *B.E. J. Theoretical Econ.* 8 (1) (Contributions).
- Glimcher, P., Dorris, M., Bayer, H., 2005. Physiological utility theory and the neuroeconomics of choice. *Games Econ. Behav.* 52, 213–256.
- Gold, J.I., Shadlen, M.N., 2001. Neural computations that underlie decisions about sensory stimuli. *Trends Cognitive Sci.* 5 (1), 10–16.
- Gold, J.I., Shadlen, M.N., 2002. Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299–308.
- Hanes, D.P., Schall, J.D., 1996. Neural control of voluntary movement initiation. *Science* 247, 427–430.
- Hebb, D., 1961. Distinctive features of learning in the higher animal. In: *Delafresnaye, J.F. (Ed.), Brain Mechanisms and Learning*. Oxford University Press, London.
- Heekeren, H.R., Marrett, S., Bandettini, P.A., Ungerleider, L.G., 2004. A general mechanism for perceptual decision-making in the human brain. *Nature* 431, 859–862.
- Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., Camerer, C.F., 2005. Neural systems responding to degrees of uncertainty in human decision-making. *Science* 310, 1680–1683.
- Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–291.
- Klyachko, V., Stevens, C., 2006. Excitatory and feed-forward inhibitory hippocampal synapses work synergistically as an adaptive filter of natural spike trains. *PLoS Biol.* 4 (7), 1187–1200.
- Körding, K.P., Wolpert, D.M., 2004. Bayesian integration in sensorimotor learning. *Nature* 427, 244–247.
- Laughlin, S.B., van Steveninck, R.R., Anderson, J.C., 1998. The metabolic cost of neural information. *Nature Neurosci.* 1, 36–41.
- Lopes, L., 1987. Between hope and fear: The psychology of risk. *Adv. Exper. Social Psychol.* 20, 255–295.
- Lord, C.G., Ross, L., Lepper, M.R., 1979. Biased assimilation and attitude polarization: The effect of prior theories on subsequently considered evidence. *J. Person. Social Psych.* 37 (11), 2098–2109.
- Ma, W.J., Beck, J.M., Latham, P.E., Pouget, A., 2006. Bayesian inference with probabilistic population codes. *Nature Neurosci.* 9 (11), 1432–1438.
- Mackworth, J., 1970. *Vigilance and Attention: a Signal Detection Approach*. Penguin Books, Harmondsworth, England.
- McKenzie, R., Turner, J., Zak, P.J., 2009. The neuroeconomics of rational decision making. In: *McKenzie, R. (Ed.), Predictably Rational? In Search of Defenses for Rational Behavior in Economics*. Springer, Berlin.
- Meyer, M., 1991. Learning from coarse information: Biased contests and career profiles. *Rev. Econ. Stud.* 58, 15–41.
- Mogenson, G.J., 1987. Limbic motor integration. In: *Epstein, A., Morrison, E.R. (Eds.), Progress in Psychobiology and Physiological Psychology*. Academic Press, New York, pp. 117–170.
- Newsome, W.T., Britten, K.H., Movshon, J.A., 1989. Neuronal correlates of a perceptual decision. *Nature* 341, 52–54.
- Nichols, J., Newsome, W., 2002. Middle temporal visual microstimulation influences veridical judgments of motion direction. *J. Neurosci.* 22 (21), 9530–9540.
- Nickerson, R.S., 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Rev. General Psych.* 2 (2), 175–220.
- Padoa-Schioppa, C., Assad, J.A., 2006. Neurons in orbitofrontal cortex encode economic value. *Nature* 441, 223–226.
- Padoa-Schioppa, C., Assad, J.A., 2008. The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nature Neurosci.* 11 (1), 95–102.
- Payne, J., 1982. Contingent decision behavior. *Psychol. Bull.* 92, 382–402.
- Payne, J., Bettman, J., Johnson, E., 1988. Adaptive strategy selection in decision making. *J. Exper. Psychol. Learn., Mem. Cognition* 14 (3), 534–552.
- Pennington, N., Hastie, R., 1993. The story model for juror decision-making. In: *Hastie, R. (Ed.), Inside the Juror: The Psychology of Juror Decision-Making*. Cambridge University Press, New York, pp. 192–221.
- Platt, M.L., Glimcher, P.W., 1999. Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Plous, S., 1991. Biases in the assimilation of technological breakdowns: Do accidents make us safer. *J. Appl. Psychology* 21, 1058–1082.
- Prelec, D., 1998. The probability weighting function. *Econometrica* 66, 497–527.
- Quiggin, J., 1982. A theory of anticipated utility. *J. Econ. Behav. Organ.* 3 (4), 323–343.
- Rabin, M., Schrag, J., 1999. First impressions matter: A model of confirmatory bias. *Quart. J. Econ.* 114 (1), 37–82.

- Ratcliff, R., McKoon, G., 2008. The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computat.* 20 (4), 873–922.
- Robson, A.J., 2001a. The biological basis of economic behavior. *J. Econ. Lit.* 39 (1), 11–33.
- Robson, A.J., 2001b. Why would nature give individuals utility functions? *J. Polit. Economy* 109 (4), 900–914.
- Roitman, J.D., Shadlen, M.N., 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Nature Neurosci.* 22, 9475–9489.
- Sah, R., Stiglitz, J.E., 1986. The architecture of economic systems: Hierarchies and polyarchies. *Amer. Econ. Rev.* 76, 716–727.
- Salzman, C.D., Britten, K.H., Newsome, W.T., 1990. Microstimulation in visual area MT: Effects on direction discrimination performance. *J. Neurosci.* 12, 2331–2355.
- Schwartz, O., Simoncelli, E.P., 2001. Natural signal statistics and sensory gain control. *Nature Neurosci.* 4 (8), 819–825.
- Shadlen, M.N., Britten, K.H., Newsome, W.T., Movshon, J.A., 1996. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* 16, 1486–1510.
- Shefrin, H., 2008. *A Behavioral Approach to Asset Pricing*. Elsevier Academic Press, Boston.
- Shefrin, H., Thaler, R., 1988. The behavioral life-cycle hypothesis. *Econ. Inquiry* 26 (4), 609–643.
- Simon, H.A., 1956. Rational choice and the structure of the environment. *Psych. Rev.* 63 (2), 129–138.
- Simoncelli, E.P., 2003. Vision and the statistics of the visual environment. *Curr. Opin. Neurobiol.* 13, 144–149.
- Spitzer, M., Fischbacher, U., Herrnberger, B., Grön, G., Fehr, E., 2007. The neural signature of social norm compliance. *Neuron* 56, 185–196.
- Stevens, C., 2003. Neurotransmitter release at central synapses. *Neuron* 40, 381–388.
- Thaler, R.H., Shefrin, H.M., 1981. An economic theory of self-control. *J. Polit. Economy* 89, 392–406.
- Timmermans, D., 1993. The impact of task complexity on information use in multi-attribute decision making. *J. Behav. Decision Making* 6, 95–111.
- Tversky, A., 1972. Elimination by aspects: A theory of choice. *Psych. Rev.* 79 (4), 281–299.
- Wickens, T., 2002. *Elementary Signal Detection Theory*. Oxford University Press, New York.
- Wolfe, J., Horowitz, T., Van Wert, M., Kenner, N., Place, S., Kibbi, N., 2007. Low target prevalence is a stubborn source of errors in visual search tasks. *J. Exper. Psych. General* 136 (4), 623–638.
- Zak, P.J., 2004. Neuroeconomics. *Philos. Trans. Roy. Soc. B* 359 (1451), 1737–1748.
- Zak, P.J., Kurzban, R., Matzner, W.T., 2004. The neurobiology of trust. *Ann. New York Acad. Sci.* 1032, 224–227.