Neurodynamical theory of decision confidence

Andrea Insabato

TESI DOCTORAL UPF / 2014

Directores de la tesi

Prof. Dr. Gustavo Deco,
Dr. Mario Pannunzi
Department of Information and Communication Technologies

UNIVERSITAT
POMPEU FABRA

The court's PhD was appointed by the recto of the Universitat Pompeu Fabra on ...., 2010.

Chairman

Member

Member

Member

Secretary

The doctoral defense was held on ....,
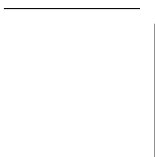
2010, at the Universitat Pompeu Fabra and scored as ....

PRESIDENT   MEMBERS

SECRETARY
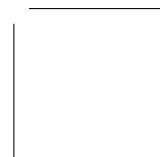
To doubt everything or to believe everything are two equally convenient solutions; both dispense with the necessity of reflection.
Jules Henri Poincaré, *Science and Hypothesis*

Acknowledgements

I believe that all the people that deserve my acknowledgment already know it. Nonetheless it seems customary to make explicit here these acknowledgments. So ... here it goes.

I guess the first thank you should go to the person that were directly involved in the work behind this dissertation.

My first (in a chronological order) supervisor Gustavo Deco deserves all my gratitude especially for introducing me in this discipline and giving me the means to progress. También le debo mis agradecimientos por los expectaculares asados que ha organizado (y por los que espero organizará).

My second (only in chronological order) supervisor Mario Pannunzi has been a guide in all these years. I deserve it to him if I have had some flash of understanding of what science is. Grazie Mario!

Results presented in this thesis (in chap. 3) were the result of a collaboration with my colleague Marina Martinez. I have to thank her for all the discussions and the work done together.

My most sincere acknowledgments go to Edmund T. Rolls for helping me in the development of the study reported in chap. 2.

I want to thank also Carlos Acuña and Jose Pardo Vazquez for recording the data used in chap. 4 and for the helpful discussions.

My acknowledgments also to the members of my PhD committee: Pascal Mamassian, Ruben Moreno-Bote and Alex Roxin, as well as to their substitutes Albert Compte and Ernest Montbrió.

To all my colleagues and friends. . . Don't ask me a personalized acknowledgment for each of you. . . I would surely forget someone (yes, I'm getting older) and that's not nice, so: Thank you all!

Grazie anche ai miei cognati e suoceri, Peppe, Bri, Gennaro e Luisa!

To my parents a thank you would never be enough: grazie per avermi "sponsorizzato" negli ultimi 28 anni!

A mio fratello Roberto e al mio fraterno amico Alveno. . . ho troppe cose per cui ringraziarvi, in parte lo sapete ed in parte no, ma comunque siete insostituibili!

Ai nonni, Livio e Feride, come non ringraziarli: siete sempre stati e sempre sarete una presenza esemplare.

Last but not least I want to thank my wonderful wife, Roberta, for supporting me during all these years and for giving me a shoulder massage while I'm writing these acknowledgments (and for dictating me this: I would have never known how to thank her for all what she does for me).

Abstract

Decision confidence offers a window on introspection and onto the evaluation mechanisms associated with decision-making. Nonetheless we do not have yet a thorough understanding of its neurophysiological and computational substrate. There are mainly two experimental paradigms to measure decision confidence in animals: post-decision wagering and uncertain option. In this thesis we explore and try to shed light on the computational mechanisms underlying confidence based decision-making in both experimental paradigms. We propose that a double-layer attractor neural network can account for neural recordings and behavior of rats in a post-decision wagering experiment. In this model a decision-making layer takes the perceptual decision and a separate confidence layer monitors the activity of the decision-making layer and makes a judgment about the confidence in the decision. Moreover we test the prediction of the model by analyizing neuronal data from monkeys performing a decision-making task. We show the existence of neurons in ventral Premotor cortex that encode decision confidence. We also found that both a continuous and discrete encoding of decision confidence are present in the primate brain. In particular we show that different neurons encode confidence through three different mechanisms: 1. Switch time coding, 2. rate coding and 3. binary coding. Furthermore we propose a multiple-choice attractor network model in order to account for uncertain option tasks. In this model the confidence emerges from the stochastic dynamics of decision neurons, thus making a separate monitoring network (like in the model of the post-decision wagering task) unnecessary. The model explains the behavioral and neural data recorded in monkeys lateral intraparietal area as a result of the multistable dynamics of the attractor network, whereby it is possible to make several testable predictions. The rich neurophysiological representation and computational mecha-

nisms of decision confidence evidence the basis of different functional
aspects of confidence in the making of a decision.

Resumen

El estudio de la confianza en la decisión ofrece una perspectiva venta-
josa sobre los procesos de introspección y sobre los procesos de evalu-
ación de la toma de decisiones. No obstante todav'ia no tenemos un
conocimiento exhaustivo del sustrato neurofisiológico y computacional
de la confianza en la decisión. Existen principalmente dos paradig-
mas experimentales para medir la confianza en la decisión en los su-
jetos no humanos: apuesta post-decisional (*post-decision wagering*) y
opción insegura (*uncertain option*). En esta tesis tratamos de aclarar
los mecanísmos computacionales que subyacen a los procesos de toma
de decisiones y juicios de confianza en ambos paradigmas experimen-
tales. El modelo que proponemos para explicar los experimentos de
apuesta post-decisional es una red neuronal de atractores de dos capas.
En este modelo la primera capa se encarga de la toma de decisiones,
mientras la segunda capa vigila la actividad de la primera capa y toma
un juicio sobre la confianza en la decisión. Sucesivamente testeamos
la predicción de este modelo analizando la actividad de neuronas reg-
istrada en el cerebro de dos monos, mientras estos desempeñaban una
tarea de toma de decisiones. Con este análisis mostramos la existencia
de neuronas en la corteza premotora ventral que codifican la confi-
anza en la decisión. Nuestros resultados muestran también que en el
cerebro de los primates existen tanto neuronas que codifican confianza
como neuronas que la codifican de forma continua. Más en específico
mostramos que existen tres mecanismos de codificación: 1. codifi-
cación por tiempo de cambio, 2. codificación por tasa de disparo, 3.
codificación binaria. En relación a las tareas de opción insegura pro-
ponemos un modelo de red de atractores para opciones multiplas. En
este modelo la confianza emerge de la dinámica estocástica de las neu-
ronas de decisión, volviéndose así innecesaria la supervisión del proceso
de toma de decisiones por parte de otra red (como en el modelo de la
tarea de apuesta post-decisional). El modelo explica los datos de com-

portamiento de los monos y los registros de la actividad de neuronas del área lateral intraparietal como efectos de la dinámica multiestable de la red de atractores. Además el modelo produce interesantes y novedosas predicciones que se podrán testear en experimentos futuros. La compleja representación neurofisiológica y los distintos mecanísmos computacionales que emergen de este trabajo sugieren distintos aspectos funcionales de la confianza en la toma de decisiones.

Contents

List of Figures

List of Tables

CHAPTER 1

Decision Confidence: An Introduction

W.C. Trow:"what is the behaviourist position on confidence?"
J.B. Watson:"I'm afraid you have come to the wrong market..."

1.1 Introduction

Decision confidence has always been considered an interesting topic of investigation since the dawning of experimental psychology (Pierce and Jastrow, 1884) and the neuroscientific community is living a renewal of interest about it in the last years due to some exciting novel findings. Decision confidence, the sensation of correctness of a choice, is an important aspect of subjective experience and a particular case of introspection (Persaud and Mcleod, 2008; Koch and Preuschoff, 2007). Moreover confidence provides an estimate of the outcome of a choice

and then proves to be very useful for planning of future actions and re-acting to a changing environment. Indeed confidence about choice and uncertainty about value estimation are important factors that influence learning and the course of action in unstable, changing environments (Rushworth and Behrens, 2008). Lack of confidence in a decision can, for example, promote a change of mind about a previous decision or can promote an exploratory strategy (Sallet and Rushworth, 2009). Therefore decision confidence is a fundamental feature of cognition giving rise to complex adaptive behavior.

Nevertheless, despite the importance of confidence, very little is known about the neural mechanisms giving rise to this feature of our cognition. This is probably due to the fact that confidence levels have always been assessed in human psychophysical experiments by means of verbal ratings that are unfeasible with animal models used in neurophysiology. In the last five years the development of new experimental procedures that measure confidence on the basis of subjects behavior opened the doors of neurophysiology to the study of confidence (Kepecs et al., 2008; Kiani and Shadlen, 2009). In parallel new models of decision confidence were proposed based on both attractor neural networks and diffusion processes (Insabato et al., 2010; Moreno-bote, 2010; Pleskac and Busemeyer, 2010) and a better understanding of the spatial and temporal construction of confidence was undertaken with psychophysical methods (Graziano and Sigman, 2009; Zylberberg et al., 2012). The main objective of this thesis is to explore and try to elucidate the neurocomputational mechanisms of decision confidence. We will work in the context of attractor neural networks composed of integrate-and-fire neurons with detailed synaptical dynamics. This framework has proved to be very fruitful in explaining many aspects of decision-making (Wang, 2002b; Marti et al., 2006; Deco and Rolls, 2006; Pannunzi et al., 2012) and its biological plausibility allow to account for (and make prediction about) neural data. Moreover some steps have been done to link these models to simpler phenomenological diffusion-like models (Wong and Wang, 2006; Wong and Huk, 2008; Roxin and Ledberg, 2008a). Therefore we think that this level of description can supply a connection between the different explanation and description levels of decision confidence.

In order to give a general picture of the phenomenon that can serve to
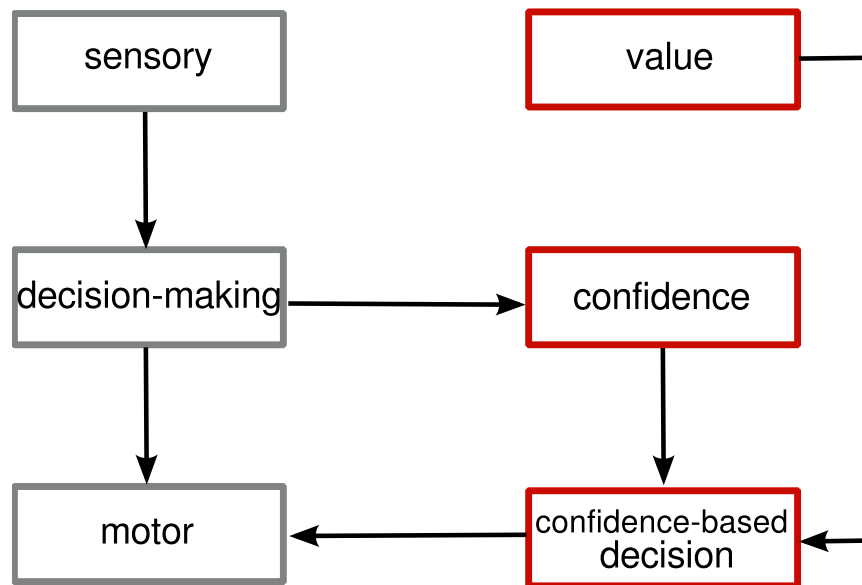
Figure 1.1: Pipeline of decision and confidence processing. Left part (in gray) represents the simplified sensory-motor path of a perceptual decision-making. On the right (in red): modules involved in confidence estimation and confidence-related decisions. The "confidence" module receives input from the "decision-making" stage, thereby implementing a sort of monitoring of the decision process. This module represents the reliability of the decision process. The "confidence-based decision" module makes a judgment based on the confidence representation coming from "confidence" module and value signals about the given options, coming from the "value" module, and transmits this second decision to the "motor" stage.

guide the discussion we would like to sketch the essential pipeline of a simple decision task involving confidence computations as in fig.1.1.

The left part of the graph represents a usual perceptual decision. In this context sensory neurons encode the relevant information about the stimulus and inform decision neurons. Hence, once the decision has been computed, the motor plan can be elaborated by the "motor" module. When the decision confidence is going to have a role in the behavioral output one needs to consider also the right part of the

graph. A new module ("confidence") can compute the confidence in the decision by monitoring the activity of the decision area. Then the confidence information can be compared with informations about the value of different options and a new confidence-related decision can be taken (e.g. the post-decision wagering experiments well summarized by Kepecs and Mainen (2012)). In this outline it is reasonable that the "confidence" module would encode continuously the decision confidence. However, if the decision confidence is ever going to have an influence on the behavior, at some point in the sensory-motor path this information need to be discretized, in order to select one course of action (e.g. in the "confidence-based decision" module for making confidence-related decisions).

Of course this is an oversimplified schema. For example, we didn't included top down influences, that could be in place at any level of the process. We also didn't take into account the role of other functional modules like the reward system or any attentional module. While the study of a more complex schema is surely valueable we wanted to propose here a very simple pipeline to first understand more closely the relationship between "decision-making", "confidence" and "confidence-based decision" modules. Indeed in this dissertation we will ask questions like: Are neurons in the brain coding the confidence on a continuous scale? Is the confidence representation abstract and task independent or is it influenced by the requirements of the environment? Are confidence neurons acting as a confidence-based decision network? To what extent can we conceptualize decision confidence as a monitoring of the decision-making process? Are the different functional modules implemented in different neural structures?

We will try to answer to these questions, more or less explicitly, throughout the next chapters.

Our discussion will be developed as follows. In this first chapter we will briefly present the state of the studies in decision-making mainly concentrating on neurophysiological findings and we will present the two main competing[1] theoretical frameworks: attractor neural net-

---

[1]We are not entering into this discussion but we want to remark that ANN and DDM are probably considered as competing frameworks more for social and historical reasons than for theoretical reasons.

work (ANN) and drift diffusion models (DDM). We will then review the main results in the study of decision confidence. We will separate this review in three sections for psychophysics, neurophysiology and theoretical models. In the next three chapters we will present the results of the investigation that brought to the writing of this thesis. In chap.2 we will present a model of Orbitofrontal Cortex (OFC) neurons that encode decision confidence (Kepecs et al., 2008). In chap.3 we will analyze the activity of neurons in the ventral Premotor Cortex (PMv) of monkeys that confirms some predictions of the model presented in chap.2 and produces new evidences about decision confidence. In chap.4 we present a new model that accounts for data recorded by Kiani and Shadlen (2009) and elucidates the mechanism of uncertain option experiments (for a discussion about the experimental procedures for confidence measuring see section 1.2.3). Finally in the last chapter a we will discuss the critical issues emerging from this work.

1.2   State of the Art

1.2.1   Neuroscience of Decision-Making

Parts of this section are included in Insabato,A. et al., The influence of spatio-temporal structure of noisy stimuli in decision-making. Submitted

Neurons encoding the various stages of a choice have been found in several brain areas using different tasks and modalities. The vast majority of these tasks, beyond the deep differences in stimulation, timing, motor output, etc. are all based on the idea of an n-alternative forced-choice (nAFC), where subjects are always required to commit to a choice between n alternatives ($n = 2, 3, ...$), even when there is no evidence at all for choosing one of the n. In this section we will review some seminal works of 2AFC, although by no means we aim to present a comprehensive review of the extant literature. In particular we will focus on perceptual decisions, leaving aside the studies on preferential choice, value based decisions, etc. In addition we will limit our

discussion to experimental results obtained in visual and somatosensory tasks, since there is a great amount of evidence about the neural correlates of decision-making in these two modalities.

In visual perception, the great richness of features of our visual experience enabled the design of a variety of decision-making tasks, including (but not limited to) the discrimination of motion (e.g. Shadlen and Newsome (2001); Gold and Shadlen (2000)), heading Heuer and Britten (2004), disparity Nienborg and Cumming (2009), and bar orientation Vzquez et al. (2000); Pardo-Vazquez et al. (2008). One of the prevalent tasks is the random dot motion (RDM) direction discrimination (e.g. Snowden et al. (1991b); Shadlen and Newsome (2001)).
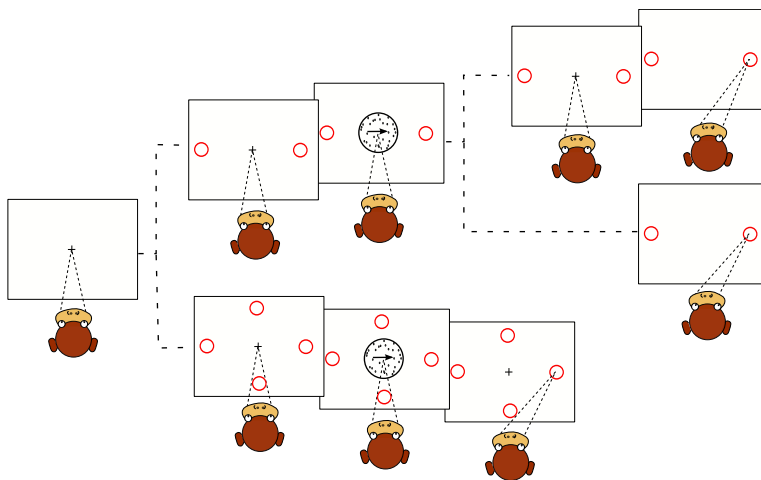
Figure 1.2: The RDM task. Top row: the fixed duration version of the task. The subject has to wait until the end of a delay interval after the extinction of the stimulus for making the saccade. Middle row: the reaction time version of the RDM task. In this task the subject can freely decide when to make the saccade. Bottom row: multiple choice version of the reaction time experiment (Churchland et al., 2008a). The subject has to take a decision between four possible directions of motion.

In this task (sketched in fig.1.2) subjects view a display where some dots have random direction of motion while others move coherently in one direction. Subjects have to decide which is the direction of coherent motion (even when there is none) and the typical response is made

by an oculomotor movement towards a visual target. The percentage of dots moving coherently determines the difficulty of the trial. This task allows to study the different parts of a decision: evidence formation, its integration into a decision signal, the holding in memory of the decision, the speed of the decision process, and the commitment to a choice. Neurons in middle temporal area (MT) are tuned to motion and therefore provide the sensory evidence for the decision Britten et al. (1992, 1993, 1996a); Shadlen et al. (1996), whereas lateral intraparietal area (LIP) and frontal eye fields (FEF) were found to integrate the evidence into a decision signal. After stimulus onset LIP neurons present a dip in firing rate and subsequently the activity differentiates according to the subject's choice: for stimuli moving towards the response field (RF) of the neuron, the firing rate ramps up, while for movements in the opposite direction, the rate decreases (see fig.1.3 for a pictorial representation). The slope of the ramping correlates with trial difficulty. Both in reaction time (RT)Roitman and Shadlen (2002b) and fixed duration experiments Shadlen and Newsome (2001); Gold and Shadlen (2000), the activity reaches an asymptotic value about 70 ms before saccade initiation, thus suggesting the existence of a decision criterion like the one postulated by diffusion-like models (see next Section).

In the somatosensory domain, the vibrotactile frequency discrimination task has also provided, in the last decades, a huge amount of evidence about decision processes.

In this task, the subject's fingertip is stimulated with a vibrator in two subsequent intervals separated by a delay (see fig.1.4). The subject must decide whether the second stimulation (f_2) has a higher or a lower frequency then the first one (f_1) and communicate the decision by pressing one of two buttons Mountcastle et al. (1990); Salinas et al. (2000). Neurons in primary somatosensory cortex (S1) have been found to increase their firing rate as a function of the stimulus frequency. Thus they encode the salient stimulus feature for the decision. During the delay between the two stimulations, the frequency of the first one must be kept in memory and neurons in second somatosensory cortex (S2), medial and ventral premotor cortices (MPC, VPC) and dorsolateral prefrontal cortex (dlPFC) were identified to encode stimulus frequency in this period Romo et al. (2002b); Brody
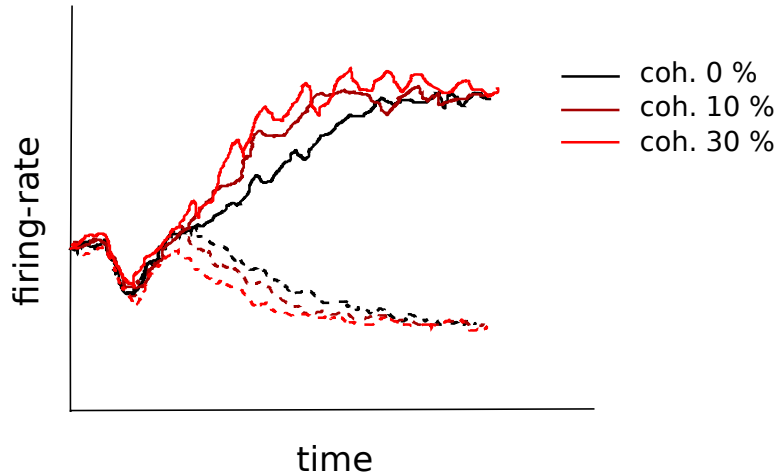
Figure 1.3: Illustration of the activity of neurons in LIP. The firing rate after the stimulus is predictive of the choice and correlates with the difficulty (the percentage of coherently moving dots).
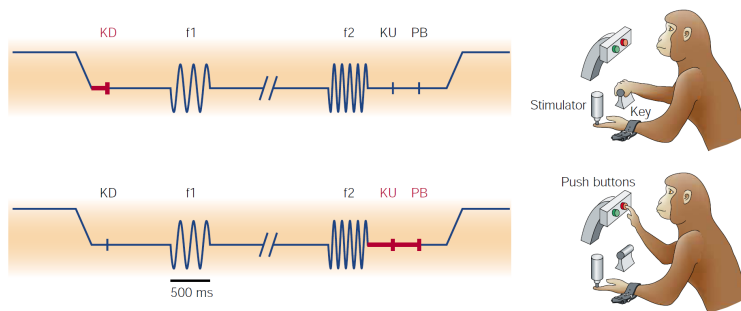
Figure 1.4: Vibrotactile frequency discrimination task. The subject holds an immovable key while a fingertip of the other is in contact with the vibrator (KD phase; top row). The fingertip is stimulated with a vibration with frequency $f1$ and after a delay a second stimulation with frequency $f2$ is applied. The subject has to compare the two frequencies and decide whether the second one was higher of lower that the first one. When the second stimulus ends the subject can release the key (KU) and communicate the decision by pressing one of two buttons (PB; second row).

et al. (2003b); Hernndez et al. (2002). When the second stimulation is applied, the comparison between $f1$ and $f2$ must be evaluated and

neurons in premotor and prefrontal cortices (and to a minor extent also in S2) encode this comparison in their firing rate, while other neurons encode either f_1 or f_2. By adding a delay between f_2 and the response, Lemus et al. Lemus et al. (2007) found that the firing of some MPC neurons during this period reflects the comparison process, while other neurons still encode either f_1 or f_2, thus suggesting a possible role for this area in the post-decision processing of the choice, as already observed in other areas Kepecs et al. (2008); Kiani and Shadlen (2009); Pardo-Vazquez et al. (2008).

While the analysis of the functioning neural systems can provide a great insight into the biological substrate of behaviour, as demonstrated by the results with neuronal recordings in monkeys discussed above, our understanding of these phenomena may also pass through the possibility of directly influence the behaviour by acting on neural systems. Following this approach, several studies have demonstrated that electrical micro-stimulation of areas involved in decision-making, both in the somatosensory [26] and in the visual [2729] domain, show similar effects to those observed when the sensory organs receive the stimulation.

Although the study of 2AFC has paved the way into the basic principles underlying decision-making, these tasks neglect important aspects inherent to most decisions that can otherwise still be considered in highly simplified experimental scenarios such as those used in typical psychophysical or neurophysiological experiments. Such aspects include the consideration of multiple alternatives, the possibility of changing one's mind or the effect that different types of irrelevant information (e.g. noise) play on decision-making.

The study of decision-making between multiple alternatives was addressed from a psychophysical perspective already in the 50s (e.g. Hick (1952)), but only in the last few years, and in the context of a RDM task, have neurophysiological recordings become available (Churchland et al., 2008b; Bollimunta and Ditterich, 2012; Louie et al., 2011) (see Churchland and Ditterich (2012) for a review). It is worth noting that theoretical attempts to account for multiple choice decision-making had already been done over the past 40 years (Tversky and Simonson, 1993; Tversky, 1972; Roe et al., 2001; Usher and McClelland, 2001;

Bogacz et al., 2007). Indeed, a family of models known as race models (Vickers, 1970) where each target (or decision) is described by an accumulator, which is close in formulation although not mathematically equivalent to DDM (Bogacz et al., 2006), can be easily extended to multiple targets by simply adding more integrators.

Churchland et al. (2008b) also reported behavioral results and for the first time recorded neurophysiological responses in monkeys (area LIP) on a two- and four-choice direction-discrimination decision task (for a representation of the task see fig.1.2 bottom row). These results have been theoretically modelled in different studies (Beck et al., 2008; Furman and Wang, 2009; Albantakis and Deco, 2009a). On one side, Beck et al. (2008) followed a probabilistic approach with special emphasis on optimality whereas Furman and Wang (Furman and Wang, 2009) and Albantakis and Deco (Albantakis and Deco, 2009a) pursued a neurodynamical approach with an emphasis on obtaining a detailed biophysical description of the circuitry underlying decision-making.

Of special interest is the situation when multiple choices simultaneously receive evidence Niwa and Ditterich (2008a) tested human participants on a 3AFC version of the RDM task. A key aspect about their experimental setting was that a multicomponent RDM stimulus was considered, i.e. the stimulus was comprised of up to three coherent motion components instead of just one direction of coherent motion. Thus, the amount of sensory evidence for all three alternatives could be controlled for. In a subsequent study, Bollimunta and Ditterich (2012) used the same experimental paradigm with monkeys while recording neurophysiological activity from LIP and suggested that a unique variable, net motion strength (NMS), in the 3AFC task is sufficient to predict monkeys accuracy and RTs. The NMS is defined from: i) amount of information associated with the highest coherence, $c_{\text{PRO}} = c_1$, and ii) average coherence of the second (c_2) and third (c_3) components, $c_{\text{ANTI}} = \frac{c_2 + c_3}{2}$, as NMS = $c_{\text{PRO}} - c_{\text{ANTI}}$, in an attempt to encapsulate in a single signal all evidence against the dominant component. Interestingly, their results seem to challenge the class of ANN models previously described, which explain very well both behavior and neural activity in decision-making tasks. Specifically, the authors suggest that competition cannot solely be mediated by lateral inhibition and indicate that feedforward inhibition is a necessary com-

ponent of the neural circuitry underlying their data. Such conclusions are based on the fact that the firing rates of decision neurons seem to show an earlier modulation due to the information providing evidence against the target than to the information in its favour. It is worth noting, however, that in a previous work (Ditterich, 2010), different diffusion models (e.g. with/without leak, lateral/feedforward inhibition) were tested on these experimental data and it was seen that all models explained equally well the behavioral data. In particular, one of the DDMs used in this study resembled a commonly used biologically plausible ANN model with one common inhibitory pool, thus suggesting that a spiking neural network could account as well for the behavioral data. Furthermore, in contrast to the conclusions derived from the reported experimental results, the NMS is unable to predict behavioral measurements in unbalanced cases, i.e. those cases in which the inputs to all pools are allowed to vary independently and the difference between the inputs to the pools with less coherently moving dots are large (this can be easily understood when considering a largely asymmetric stimulus, e.g. 50% coherence for the strongest motion components, 49% for the second, and 1% for the weakest component).

It is also worth noting that in most studies it is considered that a decision in both the DDM and the ANN framework is made once an established threshold is reached. This leads one to ask how such a mechanism could accommodate a change of mind. Resulaj et al. (2009) addressed this question experimentally by means of a psychophysical RDM task, where human subjects had to indicate the selected choice by moving a handle towards a left or right target (fig.1.5). By using continuous hand movements, as opposed to ballistic saccades, changes of mind could (occasionally) be observed in the handle traces.

Although these findings seem to pose a challenge to ANN (given the previously established stability of the decision-attractors), Albantakis and Deco (2012) showed that the attractor picture is entirely consistent with the reported experimental data. This is the case when the system operates close to bifurcation, thus separating a state of categorical decision-making from a multi-stable region. In this region, the existence of an attractor encoding the scenario where all possible alternatives fire at a high rate, makes it difficult to reach a decision, thus facilitating changes of mind. It is remarkable that a similar dynamical
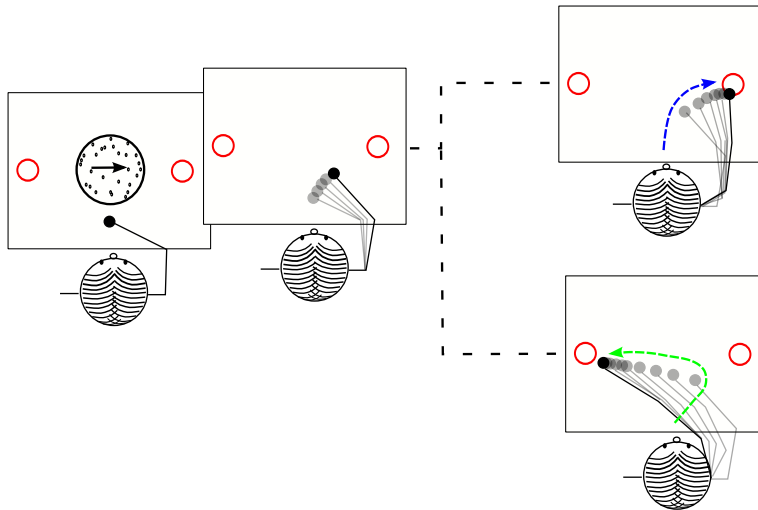
Figure 1.5: Representation of the task used by Resulaj et al. (2009) to study the change of mind in decision-making. Subjects decide about the motion direction and commit to a choice using a relatively slow hand movement. This procedure allows to record the initial preference of the subject while the ongoing integration of the evidence can still bring to a change of the initial commitment. Top row: a trial where no change occurs. Bottom row: a trial where the subject changes her mind.

regime is used here in chap.4 to account for confidence measurements in an "unsure option" task.

1.2.2 Two Theoretical Frameworks: Drift Diffusion Models and Attractor Neural Networks

As has been previously stated, 2AFC have been commonly used to investigate decision-making processes. Although several theoretical models with different flavors have been proposed, all of them share the fundamental assumption that an integration of noisy evidence over time takes place, thus accumulating such evidence until a decision is made (Bogacz et al., 2006). It is beyond our scope to describe the details of each of them, and consequently, we will only focus on the two

main competing theoretical frameworks, namely the diffusion models and the attractor neural networks.

Diffusion models

One dimensional diffusion Historically, the DDM (Stone, 1960; Laming, 1968; Ratcliff, 1978) was developed first and has been broadly used since then. The equation implementing the DDM in a 2AFC task is based on a continuous variable, $x(t)$, representing the accumulated difference between the two alternatives. In its simplest implementation, $x(t)$ is integrated over time according to:

$$dx(t) = \mu + \sigma^2 \, dW \qquad (1.1)$$

where dt is the accumulated time interval, μ is the evidence to be accumulated (inversely proportional to task difficulty and named drift rate), and $\sigma^2 \, dW$, the so-called noise-diffusion term. The value of dW is a number extracted from a normal distribution with zero mean and standard deviation equal to the square root of dt. The decision-making process is accomplished when $x(t)$ reaches one of the two boundaries: $-a/2$ or $a/2$.

The overall RT can be thought of as the sum of the time that $x(t)$ takes to reach the boundary, i.e., the decision time, and non-decision time components that account for sensory and motor processing. It is worth noting that standard implementations of the DDM may include: 1) across-trial variability in starting point ($x(0) = z$), thereby implementing the possibility of fluctuations in the starting point value from trial to trial; 2) across-trial variability in the non-decision component of processing, accounting for the possibility of fluctuations in non-decision times from trial to trial; and 3) across-trial variability in drift rate, which considers fluctuations in the drift rate from trial to trial. DDM accounts notably well for RT and performance distributions for different task procedures and speed-accuracy trade-offs (e.g. with or without time pressure; see Ratcliff and McKoon (2008) for a review). Moreover, it has been shown that DDM can reproduce the shape of the RT distributions both when it is Gaussian (Ditterich, 006a,b) and

when it has the usual positive skewness (Ratcliff and Smith, 2004; Ratcliff and McKoon, 2008). One of the main strengths of the DDM stems from the simplicity to fit its output to behavioral data (Vandekerckhove and Tuerlinckx, 2007). In this respect, DDM has been used to test a broad range of psychophysical hypotheses (e.g. Ratcliff et al. (2012)).

Race model When the task requires a decision among multiple alternatives the so called "race model" (Vickers, 1970) seems to be a more natural option. Indeed the race model represents the decision process as a race between two or more accumulators. Each accumulator is associated with a choice and the one that hits first the boundary determines the decision. The dynamics of each accumulator is governed by a diffusion process like in the one dimensional DDM. Therefore this model can naturally account for multiple choices decisions just by adding more accumulators. One dimensional DDM and race model are similar but they are not equivalent. Moreno-bote (2010) used a very clear formalism that allows to understand one dimensional diffusion model and race model as the two extremes of a continuum. Indeed we can write the equation for two integrators as:

$$dx_1(t) = \mu_1 + \sigma^2[\sqrt{1-\rho}\,dW_1 + \sqrt{\rho}\,dW_c] \qquad (1.2)$$

$$dx_2(t) = \mu_2 + \sigma^2[\sqrt{1-\rho}\,dW_2 + \sqrt{\rho}\,\nu\,dW_c] \qquad (1.3)$$

where ρ is a correlation coefficient that controls the degree of correlation between the input of the two integrators and $\nu \in -1, 1$ determines the sign of the correlation . When the $\rho = 0$ the two accumulators are independent and the system implements a race mechanism. On the contrary when $\rho = 1$ and $\nu = -1$ the two integrators are perfectly anticorrelated (each one is the antineurons of the other) and they implement a classical DDM. Fig.1.6 shows this schema.

This schema however does not take into account interactions between the accumulators (for a discussion see Bogacz et al. (2007).

As previously noted, a number of features (e.g. the average and instantaneous drift, or a change in boundaries (Ditterich, 006a,b), among
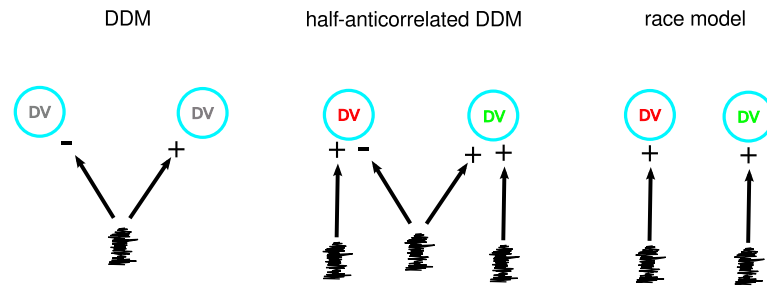
Figure 1.6: The degree of correlation of the input to the integrators control the behavior of the system. At the extremes of the spectrum there are the classical diffusion model (perfectly anticorrelated) and the race model (independent). Between the two lies the entire spectrum of half-anticorrelated models (modified from Moreno-bote (2010).

others) can be easily added to the simple versions of DDMs, thereby leading to improvements in their capability to reproduce behavioral data. However, it is still not clear which fundamental insights can be extracted from such accurate behavioral accounts. In a way, although adding and tuning new parameters may lead to substantial fitting improvements, it is not always the case that this goes hand to hand with an enhanced understanding of the fundamental underlying processes. Furthermore, one should be specially cautious when interpreting the results associated with the exploitation of the DDM fitting capabilities. Indeed a nave interpretation of Occam's razor together with an over-fitting analysis could lead to mistaken interpretations. As pointed out by Sober (1994) neither the simplicity nor the goodness of fit should be sharp criteria for choosing among different models of a phenomenon. Rather models should be selected according to their ability to survive in the experimental arena. A model should challenge new experiments with clear predictions and should be considered invalid when these predictions don't meet the experimental results.

Attractor neural networks

The DDM is a phenomenological model, and therefore, it does not attempt to provide a detailed description of the neural mechanisms

underlying decision-making. In contrast, nonlinear ANN models of spiking neurons crucially seek a biophysically inspired description of the processes underlying decision-making. These type of ANN had been initially used to explain the neurophysiological basis of other cognitive functions such as working memory (Amit, 1989; Amit et al., 1994; Brunel and Wang, 2001a). Indeed, the observation that besides decision-related activity LIP neurons also exhibited persistent activity during delay periods (Shadlen and Newsome, 2001) inspired Wang to explore the possibility that ANN of working memory could also explain the integration of stimuli and the formation of perceptual choices Wang (2002a).

The basic computational units of these ANN models are neurons represented usually with the integrate-and-fire model. Here we briefly describe the system while all the mathematical details are given in the appendix A. The integrate-and-fire model describes the membrane voltage of a neuron through a differential equation until a voltage threshold is reached. The crossing of the threshold triggers a spike which is taken to be a stereotyped event - i.e. when the threshold is crossed the spike count is incremented by one, the membrane potential is reset to a predetermined value and a refractory period follows in which the neuron doesn't integrate the input. Usually the incoming input spikes to a neuron are processed through dynamic synapses that represent different ionic channels: α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) and N-Methyl-D-aspartic (NMDA) for excitatory connections and γ-Aminobutyrric acid (GABA) for inhibitory connections. Nonetheless sometimes instantaneous synapses are used in conjunction with synaptic delays (see e.g. **?**, p.125-153 and **?**).

Typical configuration of ANNs used to account for decision-making processes will be organized into $n + 2$ populations (pools) of leaky integrate-and-fire neurons with common inputs and connectivities, where n corresponds to the number of choices in nAFC tasks. The n integrators are implemented by pools of excitatory neurons that respond selectively to evidence in favor of one of the possible decision targets. Moreover, a homogeneous pool of inhibitory neurons, globally connected to all neurons in the network, and a pool of excitatory neurons, which is not selective to any of the directions of motion, are

also considered (see fig.1.7). The models have an all-to-all connectivity with recurrent connections between cells from the same selective pool increased by a factor $\omega_+ > 1$ with respect to the baseline connectivity level, and weakened connectivities by a factor $\omega_- < 1$ between cells from different selective pools, following the hypothesis of Hebbian plasticity (i.e. synaptic efficacies are strengthen when pre- and post-synaptic neurons activities are correlated). This structure of synaptic weights is a key aspect in the formalism of attractor dynamics, which endows the system with the capability to implement a biased competition of the different populations of excitatory neurons that is mediated by inhibition, and establishes the competition and cooperation processes as the basic elements of the underlying neural computations. Therefore ANNs model the decision process as a competition between neural pools biased by the evidence for the decision. This system differs from a race model in that the race does not implement interactions between the "competitors" (however we will see later that a diffusion process is a valid reduction of ANN in a particular condition).

In ANN models, the long-term behavior of the nonlinear dynamical system defined by neural networks of interconnected neurons, is described by the so-called *fixed points* that partition the configuration space into basins of attractions. Such basins arise from the initial configurations of the system, which lead to the same attractor. Fig. 1.8 illustrates an example landscape of attraction basins. In this theoretical framework, 2AFC decision-making can be modeled by an attractor network with a minimum of two stable fixed points, which represent the two alternatives. Such a system would display bistability and the transition from an initial configuration towards one of the two stable attractors (i.e. stable unless a sufficiently large amount of noise takes the network out of the attractor) would correspond to the decision process. In these models, the usual way to account for RTs, a decision is considered to be made whenever the activity of one of the pools reaches a threshold (but see the last paragraph of sec.1.2.1 for a brief discussion).

The parameters of the ANN can shape the attractors landscape in different ways. The bifurcation diagram is a useful representation that shows the changes in attractors landscape due to the modification of a parameter. Fig.1.9 (modified from Deco et al. (2013)) shows an
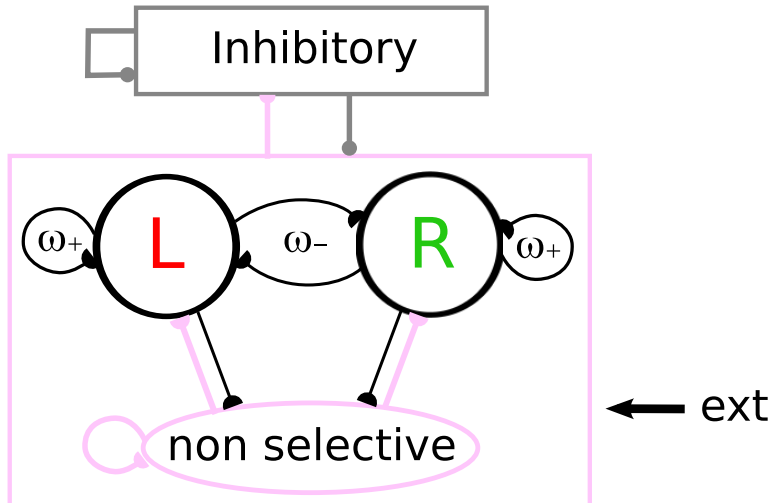
Figure 1.7: Architecture of an example ANN network for 2AFC decision-making. The connectivity in the network is all-to-all and strength of connections is assigned according to a Hebbian rule. The selective pools R and L inhibit each other indirectly through connections to the inhibitory pool and diminished mutual excitatory connections (ω_-). The increased recurrent connectivity of selective pools (ω_+) and the mutual inhibition produce a competition between selective pools. The external input can bias the competition in favor of one or the other pool.

example bifurcation diagram obtained varying the input strength (frequency of the incoming external spike train) to the selective pools. This parameter is important in ANN since it can manipulate the speed-accuracy trade-off of the system. Indeed we can see that for very low values of the input only the spontaneous attractor is stable and hence no decision can be taken. After the first bifurcation two decision attractors appear. Depending on the stability of the spontaneous symmetric state the system will work in a multistable or bistable regime exhibiting different behaviors (an example of this difference is evident in the results of chap.4). Finally for very high intensity of the input the decision attractors disappear and again no decision is possible.

Especially in the last decade, as experimental evidence grows, both competing theoretical frameworks succeed to account for the reported
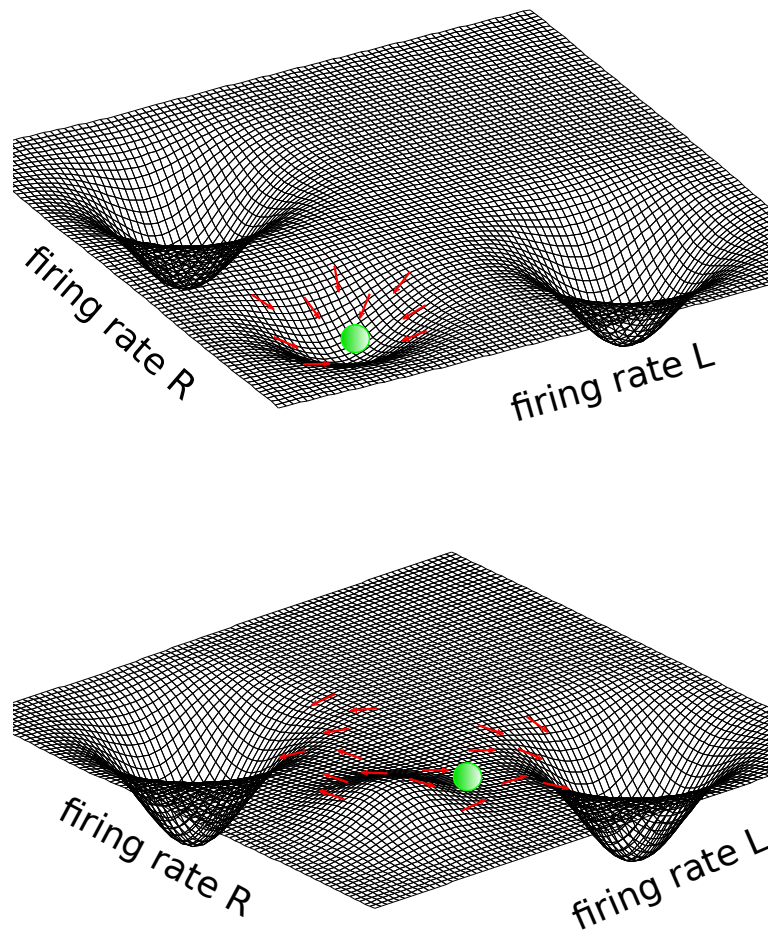
Figure 1.8: The landscape of an example ANN. The dynamical system defined by the neural network can be understood with the analogy of a ball falling on a surface. The ball will fall along the direction of local maximal gradient. The wells represent the fixed points of the system since at the bottom of the well the potential energy of the ball is at a minimum and it want escape from that state. In the first configuration (left) the spontaneous state (with both pools having low firing rate) is stable and the system remains there until noise fluctuations bring it into the basin of attraction of another fixed point. In the second configuration (right) the spontaneous state is no longer stable and the system dynamics evolve towards one of the two decision attractors (note that this is only an illustration and not the only possible configuration of a decision mechanism).
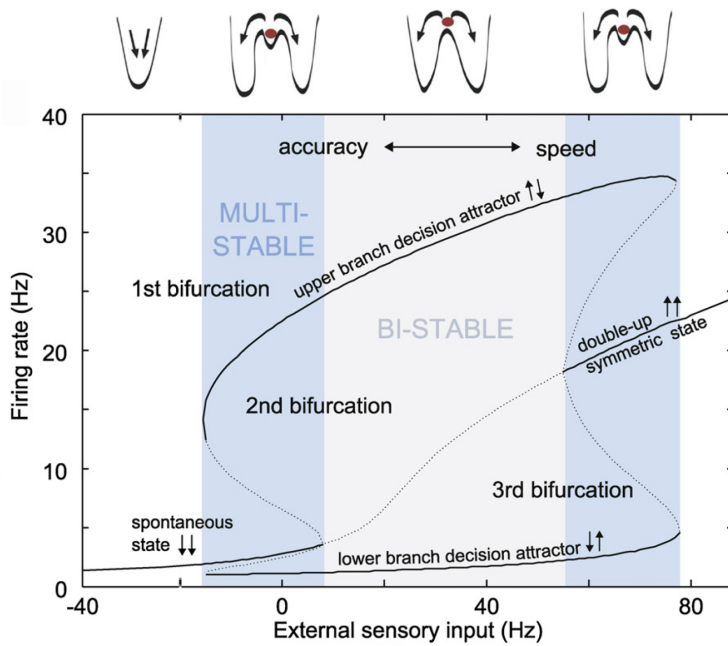
Figure 1.9: Bifurcation diagram of an example ANN. The input strength modifies the attractors landscape producing in turn different trade-offs between speed and accuracy of the decision. Solid lines indicate the firing rate of selective pools for stable solutions of the dynamical system (in the spontaneous and symmetric states both pools have the same firing rate). Dotted lines shows unstable solutions. The illustration on the top shows the hypothetical 2D potential profiles associated with the different dynamical regimes. Modified from Deco et al. (2013).

findings. One such example can be illustrated by the observation that motion pulses influence both behavior and LIP neural activity, with the later pulses being less relevant than earlier ones (Kiani and Shadlen, 2005; Kiani et al., 2008). A DDM with a leakage term could reproduce this experimental finding, while the time-varying dynamics of the attractor model explained both behavioral and neural data (Wong et al., 2007; Wong and Huk, 2008) as much as the simplest DDM (Kiani et al., 2008).

At the expense of a poorer biological plausibility, one of the great

advantages of DDM over the ANN is the fact that the DDM is described by a single equation. In contrast, ANN are endowed with richer dynamics, thus allowing to model neurophysiological data (i.e. neuronal spiking activity) that can be subsequently used to derive behavior. Moreover, the mean-field approach (Brunel and Wang, 2001a) can also reduce the amount of equations of the ANN, thus leading to a formal framework that allows to treat the dynamical system analytically. With this approach, the number of equations is proportional to the number of different populations of neurons. Notably, a further step has been done with an approach that combines numerical and analytical methods (i.e. mean-field) to reduce the system to two rate equations in Wong and Wang (2006). Later, Roxin and Ledberg (2008b) derived a formal relationship between the mean-field reduction of the ANN and a one-dimensional nonlinear diffusion in the proximity of the bifurcation to bistability, where the spontaneous state destabilizes. This is a valid reduction for all winner-take-all models, and allows to relate the variables of the nonlinear diffusion process to those of the full spiking-neuron model, and thus, to neurobiologically meaningful quantities.

Regarding the biological plausibility of the DDM another approach, different from the one of Roxin and Ledberg (2008a), has been taken by Smith (2010). He first shows that a Wiener process is equivalent in the long time to an integrated OU process. As already well known from the Stein model Stein (1965), an integrated OU process can be approximated to a pair of opponent shot noise processes (when their intensity is very high). Thus, the link with neurophysiology can be established in that shot noise processes have been used to model neural responses to action potentials. This approach is interesting and independent of the ANN formulation bu nonetheless, we note that most biologically plausible models of decision-making are not based on single neuron responses but rather on population dynamics in structured networks. In a subsequent study, Smith and McKenzie (2011) provide an alternative analysis that demonstrates how a time inhomogeneous OU velocity process emerges even in the context of a simple recurrent architecture. These works are not conclusive and further analyses are needed to shed light on the relations between ANN and DDM and on the other possible neural implementations of DDM.

Mechanism for the commitment to a choice

An issue, which is less considered and that is common to Both the diffusion and the ANN framework is that of the choice mechanism. By choice mechanism we mean here the way a commitment to a choice is determined in decision-making models.

DDMs decision variable keeps diffusing until it reaches the absorbing boundary (if this parameter is not set to infinity). Although the stationary condition induced by the boundary is produced externally, this state could be regarded as a decision state. In order to read out this decision state, historically, DDMs used a fixed threshold (Ratcliff, 1978). This mechanism is compatible with neurophysiological findings in LIP, as was already explained above. And yet, when facing fixed-time experiments, some investigators disregard the threshold and determine the choice based on the sign of the decision variable alone (e.g., Kiani and Shadlen (2009); Brunton et al. (2013)).

In ANN models the decision is given by the position of the system in the attractors landscape. Even in the condition of no evidence for the decision an ANN will reach a decision attractor and stay there (although change of mind are possible as shown by Albantakis and Deco (2012)). Therefore we could say that the ANN reaches a decision state whenever it enters into the attractor. However this condition is intrinsic to the decision network and it has to be read out by another network in order to produce a motor plan. The mechanism to read out the decision state is what we refer to as choice mechanism. In 2AFC tasks, since the two decision attractors are separated in the 2D space defined by the firing rates of the decision pools, different possible choice mechanisms can be used. The most frequently used is a threshold on the activity of the decision pools (resembling the classical DDM choice mechanism), but a mechanism based on the difference of activity between pools is also sometimes used (Marti et al., 2006; Pannunzi et al., 2012). When considering multiple alternatives, several functions of the state of the integrators could be used (e.g., difference between the two larger accumulators, between extremes, between the largest and the mean of the others, etc.). However more research is necessary to further constrain the models. The experimental paradigm proposed by Niwa and Ditterich (2008b) whereby different amounts of

evidence can be provided to each of the components seems an ideal candidate to shed some light on this issue.

1.2.3 Neuroscience of Decision Confidence

In this section the main neuroscientific findings about decision confidence will be presented. The first part is devoted to important psychophysical results. The second part explains the recent neurophysiological evidence for the neural basis of confidence and the last part gives a brief overview of theoretical accounts of decision confidence.

The problem of measuring confidence

Different experimental paradigms have been used to measure the confidence in a decision. Since the dawning of experimental psychology (Pierce and Jastrow, 1884) verbal ratings were largely used with humans subjects. In particular confidence ratings were used to reconstruct the experimental receiver operating characteristic (ROC) curve in the framework of signal detection theory (SDT) (Green et al., 1966), but there is no homogeneity in the scale adopted for the rating: Garret (1922) used a percentage scale, Foley (1959) employed the words "suppose", "think", "sure", "certain", "positive", Green et al. (1966) used a six point numerical scale, Henmon (1911) used letters from "a" (confident) to "d" (doubtful), Watson et al. (1964) and recently Graziano and Sigman (2009) used a continuous scale and a sliding pointer. It seems anyway that no difference was found between discrete and continuous scales (Rockette et al., 1992). Indeed Rockette et al. (1992) compared the confidence judgments on a five category and a continuous scale by radiologists about the presence of a mass in an abdominal computer tomography. They report no significant differences in the accuracy of confidence judgments as detected by an ROC analysis.

Animals are not able to give verbal report about their confidence, therefore other methods have been developed to measure it. This methods can be roughly classified in two types: uncertain option tasks and post-decision wagering tasks (for a good review see (Kepecs and

Mainen, 2012); fig.1.10 summarizes the different ways of measuring confidence).

rating post-decision wagering uncertain option
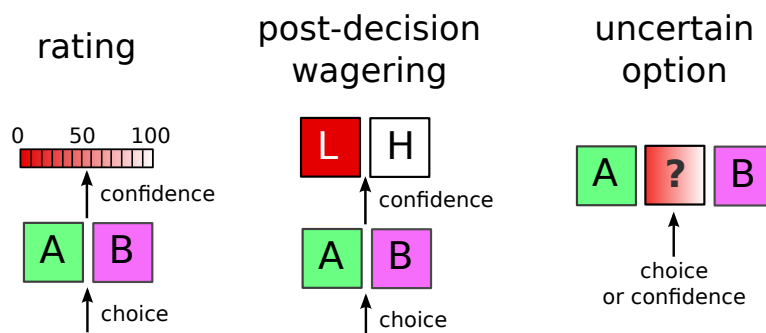
Figure 1.10: Schema of the principal three confidence measures (adapted from Kepecs and Mainen (2012). The confidence rating is the usual measure emploied with human subjects. It allows to record both the decision and the confidence judgment (where the scale of confidence can be discrete or continuous). The post-decision wagering can be used also with non human animals and also allows to record decision and confindence. However the number of confidence category is usually binary and, although it would be possible to design a continuous bet paradigm, to our knowledge, it has been not done yet. The uncertain option task is also feasible with animals and allows the recording of either the choice or the confidence level on each trial: If the subject choose the target associated with the sure reward a low confidence in the decision is implied but the chosen option (between A and B) can not be known.

Experimental paradigm using an uncertain option are actually not new in experimental psychology (Angell, 1907; Watson et al., 1973) but in the last twenty years they came into the focus of attention as a way for rigorously study confidence in animals and humans (Shields et al., 1997; Smith et al., 1995, 1997, 2003). In this task a stimulus is presented that drives a binary decision or classification (e.g. a pattern of random dots moving in one of two directions, or a tone that needs to be classified according to a given reference threshold). Subjects have not two but three possible choices, e.g. left option, right option and an "uncertain" option. The idea is that subjects would choose the uncertain option when they lack confidence in the perceptual judgment (as confirmed by post-experimental questionnaires (Smith et al.,

1995)). Smith and colleagues compared the results of humans and animals (dolphins, monkeys, rats) in this task and found that the dolphins, monkeys and humans had similar response distributions. This task however is only a weak prove of metacognitive ability and it could be not appropriate to measure confidence since the uncertain option could be simply associated with features of the stimulus intermediate respect to the two extreme categories, reducing to a mere multiple choice decision-making task. A partial solution to this problem was given by a similar task employed by Hampton (2001b). In this experiment monkeys could decide whether they wanted to answer to a memory test and get a preferred food or to decline it and receive a non-preferred food. The structure of the task is similar to that of Smith's experiments and indeed, since the difficulty of the task was simply manipulated by varying the delay between stimulus and response, the subjects needed only to associate the decline option with longer delays. However the probability of correct in forced choice trials (when the decline option was not available) was lower than that in free choice trials. This means that subjects can access information about the expected outcome of the judgment and hence opting for the decline option denotes low confidence in the decision. Nonetheless an alternative explanation of the increased performances is possible, which weaken the link of this behavior with confidence. Indeed fluctuations in the general vigilance state of the subject could also produce a higher probability of correct in free choice trials (since subjects would accept the perceptual task only when they have high vigilance). Although this would still imply a metacognitive process it would be different from a confidence judgment. We are going to address this problem from a computational perspective in chap.4.

Another weakness of the uncertain option task is that it allows to record either the response only or only the confidence in each trial. On the other hand post-decision wagering paradigms (Persaud and Mcleod, 2008) allow to obtain both informations. In a post-decision wagering task, after deciding, the subject has to bet about the correctness of her choice. It is expected that subjects bet higher in confident respect to uncertain trials. Kepecs et al. (2008) adapted the post-decision wagering task in order to use rats as subjects. They delayed the feedback after the choice allowing the animals to initiate a new

trial instead of waiting for the reward. Using this task they found, in contrast to Smith et al. (1995), that rats behavior shows the hallmark of confidence. However a limitation of their paradigm was that the confidence bet was only binary and therefore only a single bit of confidence information could be gained on each trial. Shields et al. (2005) tried to use a form of post-decision wagering with monkeys in order to mimic confidence reports in humans. Unfortunately they found that monkeys responses were similar to that of human subjects only for two category wagering (high *versus* low confidence). When using three options to bet about the perceptual choice monkeys behaviour was different from that of humans. In order to improve the post-decision wagering task and obtain a graded measure of confidence on each trial Kepecs and Mainen (2012) present a variation of the task, where the delay between choice and feedback is quite long and sampled from an exponential distribution. The time that the subjects are willing to wait for the reward is an indicator of the confidence that they have in the decision. Indeed the authors report that the waiting time present the characteristic pattern of confidence (described below).

Psychophysics of decision confidence

In this section we will highlight the main relations between confidence and several variables that emerged in many different psychophysics experiments (however we only review, in general, studies based on perceptual decision tasks). The variable that we take into account in the following are: discriminability, reaction time, accuracy, speed-accuracy trade-off (SAT), expectation.

Discriminability

The first studies about confidence put this variable in relation with the discriminability of the stimuli used in the experiment. In a task of lifted weights, where the subject has to distinguish the greater of to weights, Garret (1922) found confidence to be a monotonically increasing function of the difference between the two weights. These results were confirmed by Johnson (1939) and Festinger (1943). Both studies used a two-category discrimination task, in which the subjects were required to indicate the longer of two lines. They found that confidence

ratings plotted against the difference between the stimuli resembled a sigmoid like the classical psychometric functions for accuracy. These early results were confirmed by the subsequent research also using animals both with the uncertain option task (Kiani and Shadlen, 2009) and with the post-decision wagering task (Shields et al., 2005; Kepecs et al., 2008). These studies can be represented through signal detection theory as the comparison of two random variables. For example, in the experiment of lifted weights we can immagine that on each trial the perceived weight of the first object is a sample of a distribution of values whose variability is given by the noise in the sensory system; similarly the perceived weight of the second object will be another random variable. The decision about which weight is bigger involve a comparison between the two random variables. Usually (e.g. in the experiments mentioned above) the discriminability is manipulated by varying the distance between the means of the distributions. Anyway other possible manipulations would imply a change in the variability while holding the mean constant and a change in both the mean and the variability. However, to our knowledge, no results have been publish exploring these conditions. It is worth noting that also in the motion discrimination task the confidence in the decision has been found to decrease as a function of both stimulus duration and percentage of coherently moving dots (the two variables that manipulate the discriminability of the motion direction) (Kiani and Shadlen, 2009).

Up until now we only considered correct trials but if one looks at error trials the relation between confidence and difficulty is mirrored. First Pierrel and Murray (1963) founded that confidence was lower for error trials rather than for correct trials. Later it has been found that the confidence in a decision decreases monotonically in error judgments as a function of the ease of the trial (Vickers et al., 1985; Kepecs et al., 2008; Kepecs and Mainen, 2012). The confidence measure, when plotted as a function of the difference between the stimuli to be compared, produces a characteristic X pattern if correct and error trials are considered separately (see e.g. fig.1.11).

We also note that a recent paper addressed the problem of the construction of confidence using a partial report task (Graziano and Sigman, 2009). In this study the main parameter that manipulate the difficulty is the duration of the interval between the stimulus and the cue.

In this study the authors report that the confidence in the judgment decreases as a function of the duration of the interval and the same modulation was found for correct and error trials. However the task employed is quite complex and has a strong memory component and therefore conclusions based on these results should be generalized with caution.

Reaction time

Confidence was also found very early to be an inverse linear function of reaction time (Henmon, 1911; Volkmann, 1934). These early proposals were motivated also by the fact that confidence seemed to always underestimate the probability of correct responses, saturating very early (see further for a more complete description of the relation between confidence and accuracy), and therefore it should be based on some other feature of the decision process, like the time taken to commit to a choice. Later Reed (1951) stated that confidence can be modeled as $c = a/t + b$, where c is confidence, a is constant and b is the reaction time for an infinitely large stimulus difference (i.e. in a case where subject achieve 100% correct responses). Audley (1964) in an experiment where subject had to compare the relative frequency of green and red light flashes also reported an inverse relation between confidence and reaction time for any level of discriminability. Pierrel and Murray (1963) reported that in a lifted weights experiment confidence increases with stimulus difference and was lower for error trials than for correct trials. They found also that reaction times were inversely related to confidence ratings and higher on error trials than on correct trials.

Irwin et al. (1956) presented an experiment that seems in contrast with evidence presented so far. They asked subjects to inspect cards from a 500 pack with positive and negative numbers printed on them and to decide after each card if the mean of the pack was greater or less than zero. They also recorded confidence ratings and found that confidence increased with the number of observations made by the subject before making a decision. While this result seems puzzling at first there is a fundamental difference between the task of Irwin et al. (1956) and the tasks presented so far. In lifted weights or line length experiments all the evidence supporting one or the other alternative is given to the

subject in a very short time window and then the subject can take her time to evaluate that evidence. In the card experiment evidence is accumulated discretely at a very low rate, therefore evidence increases as a function of time and confidence could just be following the process of evidence formation. Moreover, to increase the complexity of these results, Audley (1964) reported an inverse relation between confidence and response time, analogously to other simple perceptual judgment experiments. In their task each flash concurred to the evidence in favor of one or the other alternative so their result seems to contradict the one of Irwin et al. (1956). However we note that if the arrival rate of evidence packets is quite high the decision problem start to resemble the line comparison task. Therefore in an experiment varying the arrival rate of information for the decision one should observe a transition from a phase where confidence increases with response time and another phase where an inverse relationship can be observed. To our knowledge such an experiment has not been done.

Accuracy

Since confidence is thought to be an estimate of the outcome of the decision it is not surprising that its relation to accuracy was studied very early. The first attempt was done already by Pierce and Jastrow (1884). They founded that confidence was a direct function of the accuracy and could be described as $c = h \cdot log(p/(1-p))$, where p is the probability of a correct response and h is a constant. It was recognized very soon that the confidence judgement is a complicated and non-linear function of the objective accuracy function. Garret (1922) reports that confidence was not a reliable prediction of objective accuracy, an observation confirmed also by other studies (Johnson, 1939; Festinger, 1943; Baranski and Petrusic, 1994). In particular Baranski and Petrusic (1994) report that subjects are overconfident in difficult decisions and underconfident in easy decisions and that both calibration and resolution of confidence decreased when the difficulty of the decision increased. Using rats Kepecs et al. (2008) found that the firing rate of neurons encoding uncertainty is a decreasing linear function of the probability of correct responses in an odour categorization task. In a similar task Kepecs and Mainen (2012) also showed that the waiting time, that they propose as a valid measure of confidence (see above section 1.2.3), is an increasing function of the accuracy. Nonetheless,

going back to human subjective reports, Green et al. (1966, p. 106) show some heterogeneity in the calibration of different subjects (with some subjects obtaining fair calibrations). Moreover, recently, Barttfeld et al. (2013) have shown that the reliability of confidence judgments present a high variability over subjects. They found that the ROC curve of type II performance (i.e. how much the confidence report predicts the objective accuracy) corresponding to 25 subjects vary a lot. The ROC curve of type II performance stays close to the diagonal (bad prediction) for some subjects but has a very bowed profile (and large area under the curve, i.e. good prediction) for other subjects.

Speed-accuracy trade-off

In the experiment of Johnson (1939), authors used different instructions ("speed", "usual" and "accuracy") to tell subjects, respectively, to be rapid, to do just as usual the task or to be as accurate as possible. They found that the different instructions elicited different behaviour and characterized by both different reaction times and probabilities of correct responses. Nonetheless the three conditions had almost no effect on confidence judgements. Some decades after Vickers and Packer (1982) report that in a similar experiment in "accuracy" trials, the probability of correct responses was higher, the RT was longer and the confidence was higher compared to "speed" trials. Further investigation in this area could bring new results that can be used to constrain the available model of confidence processing (e.g. the race model, described in sec.1.2.3, can easy account for the results of Vickers and Packer (1982) but not for that of Johnson (1939)).

Neurophysiology of decision confidence

Neurophysiological evidence about decision confidence has become available only in the last 5 years mainly thanks to two studies that used respectively rats and monkeys as animal models. A work by Kepecs and colleagues described neurons in the rat Orbio-Frontal Cortex (OFC) that encode decision confidence estimates (Kepecs et al., 2008). The behavioural paradigm consists in a binary odour categorization task. The rats had to perform the binary categorization task with a mixture of two pure odorants (A, caproic acid; B, 1-hexanol), by entering in

one of two ports to indicate that the mixture was more like odour A or more like odour B (see Fig.1.11a). Correct choices were rewarded after a variable delay of 0.3-2 s. Varying the relative concentration of the odorants allows to alter the difficulty of the trial. Neural activity related to decision confidence should occur just after the decision is taken and before the trial outcome. Kepecs et al. (2008) therefore analyzed recordings of neuronal activity during this delay period. The neurons were then divided into two groups based on whether they fired with a higher frequency on correct or on error trials. Kepecs et al. (2008) found that the group of neurons with an increased firing rate on error trials had higher firing rates with easier stimuli. The same neurons fired at a substantially lower rate on correct trials, and on these trials the firing rates were lower when the decision was made easier. This produced the typical X pattern associated with confidence (showed in Fig.1.11c,d). Hence the neural activity in OFC seems to jointly reflect the decision difficulty and the trial outcome, and thus represents decision confidence. However in this experiment there was no behavioral response about confidence, i.e. the rats took only a perceptual choice but the confidence in that decision was not recorded. Therefore one could reasonably doubt that the neural activity of OFC neuron, as recorded by Kepecs et al. (2008), is actually correlated to decision confidence. This criticism is even more relevant, given that the existence of metacognitive functions in rodents has not been proven.

In order to assess this possible criticism, authors performed a second experiment to investigate if rats were able to make use of the information encoded by OFC neurons. To this aim they used a post-decision wagering experiment. The delay period between the response and the reward was prolonged up to 8 s in order to allow the rat to reinitiate the trial. The subject could decide to leave the choice port and restart the trial, or could wait for the reward. The decision to stay or to return to odour port should be related to the level of confidence of the rat about the previous choice. Indeed if the subject is sure of its decision it should wait for the reward, but if its decision was made with low confidence, e.g. if the odour percentage is around 50%, the subject would probably reinitiate the trial hoping that the next stimulus will be easier. It was found that when the likelihood of a reward was low, due to the decision difficulty and the choice just made, the rat
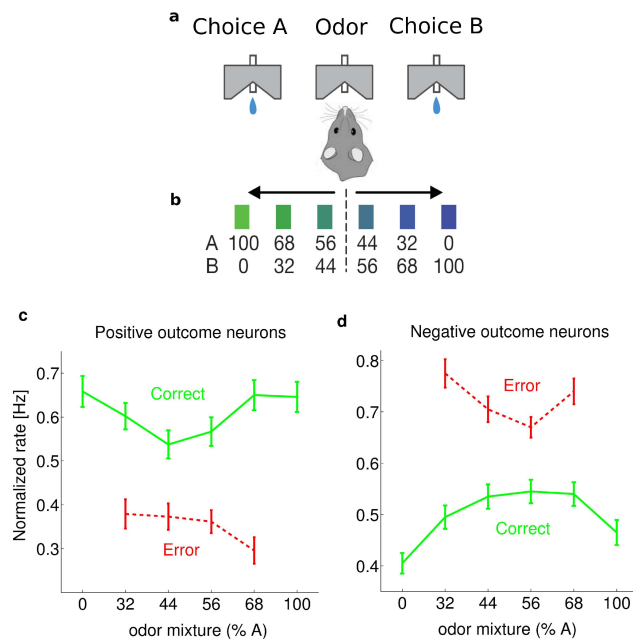
Figure 1.11: a. Odour mixture categorization task: when the rat enters the odour port, an odour mixture is delivered. The stimulus is defined by the percentage of the two pure odorants A and B in the mixture as shown in (b). Then the subject categorizes the mixture as A or B by moving left or right, according to the dominant odour. c. Mean normalized firing rate of the positive outcome selective neuronal population. These neurons fire faster on correct trials then on error trials, and increase their firing rate as the task becomes easier on correct trials, while they decrease the firing rate on error trials. d. Mean normalized firing rate of the negative outcome selective neuronal population in the orbitofrontal cortex as recorded by Kepecs et al (2008). These neurons have sustained activity on error trials, increasing their firing rate further when the decision is made easier. On correct trials they fire slower, and decrease their firing rates as the decision becomes easier. (modified from Kepecs et al., 2008).

returned to the odour port. The probability that the rat would restart a trial as function of stimulus difficulty and accuracy show the typical X pattern and reflects the responses of OFC neurons (see Kepecs et al., 2008, Fig.5). In chap.2 we will discuss further these results and

present a model that accounts for these data.

Another recent paper from Kiani and Shadlen addresses the issue of decision confidence in a different manner. Kiani and Shadlen (2009) used a modified version of the 2 alternative forced choice task (2AFC) introducing an uncertain option in order to distinguish trials in which the monkey had low confidence associated with the perceptual decision. They also recorded single neurons activity from the Lateral Intraparietal sulcus (LIP). In the usual 2AFC paradigm a random dots motion (RDM) stimulus is shown to the monkey, who has to decide to make a saccade towards one of two targets, located at the opposite sides of the visual field (T_{in}: The target in the response field (RF) of the recorded neurons; T_{opp}: The target at the opposite side), according to the prevailing motion direction. In this implementation of RDM a subset of dots is just randomly relocated at each new frame producing the illusion of random movement while another subset moves coherently in one direction. The percentage of coherently moving dots can be modified allowing control over the evidence for the decision. The unsure option was implemented as a third target, called "sure target" (T_s), shown orthogonally to the axis between the motion targets. T_s was associated with a lower but sure reward. As explained above, if the monkey is not confident about the perceptual decision, then he can choose T_s and get a reward. The probability of choosing T_s reflects the animal's average confidence in a set of trials, while on single trials either the outcome of the decision or a binary value of decision confidence can be extracted.

Psychophysical results show that confidence decreases when the fraction of coherently moving dots decreases and also when the duration of the RDM display decreases. Although the two parameter may exert different effects on the decision process it could be said that confidence decreases when the task is made more difficult (see Fig. 1.12).

Neural recordings show the well-known choice related activity of LIP neurons (Fig. 1.13c) in standard trials (see Gold and Shadlen, 2007). During trials in which T_s option was given but waived neural recordings seem unaffected by the third possible choice. On the other hand, when the subject chose the sure target the firing rate during the stimulus period remains lower than in standard trials, staying in a middle
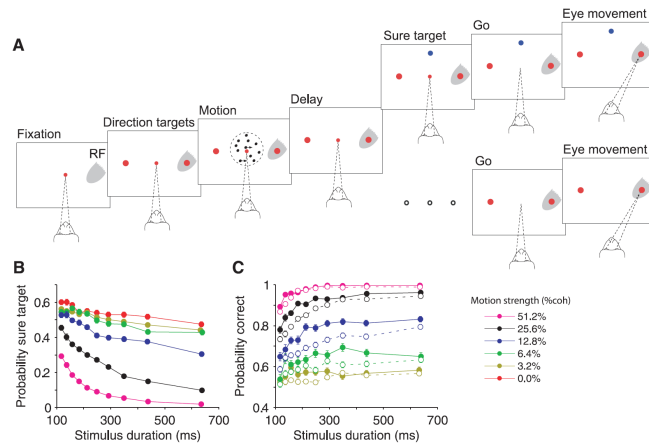
Figure 1.12: a. Experimental paradigm in trials in which T_s was shown (upper branch) and in standard trials (lower branch). b. Probability of choosing T_s as a function of stimulus duration grouped by stimulus coherence. c. Accuracy as a function of stimulus duration and grouped by stimulus coherence. Open circles represent accuracy on standard trials, while filled circles represent accuracy on trials in which T_s was shown but waived. A general improvement of performance is notable when monkey waived the sure target. (modified from Kiani and Shadlen, 2009).

level, both for trials in which prevailing motion was toward T_{in} and trials in which it was toward T_{opp}, during the stimulus period. After the stimulus period firing rates attenuate until the end of the trial (Fig. 1.13d).

Authors conclude that the activity in LIP neurons is really informative about animal's decision confidence, because when firing rates don't achieve a sufficient level (above a given threshold) confidence in the decision is too low and the sure option is chosen. We will discuss further these results in chap.4, where we present a model to account for these data.

Theoretical models of decision confidence

The psychophysical results summarized in Sec. 1.2.3 and 1.2.3 were explained using the framework of diffusion models (e.g. (Vickers, 1979a;
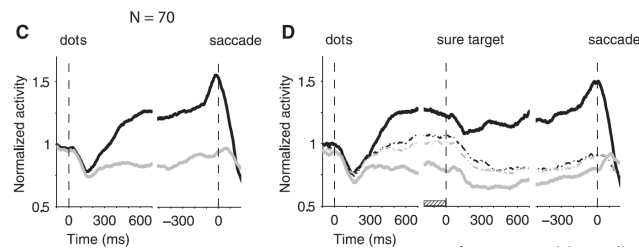
Figure 1.13: c. Average neural activity of LIP neurons over time when T_s was not presented. Black and gray traces represent respectively choice for T_{in} and T_{opp}. d. Average neural activity when T_s was shown. Solid lines and dashed lines represent trials in which the sure option was waived and chose, respectively. (modified from Kiani and Shadlen, 2009).

Kiani and Shadlen, 2009)).

Vickers (1979a) used a race model to account for several features of decision-making processes and suggested that confidence can be encoded in such a model in the difference between the two decision variable at the time of the choice, what he called the "balance of evidence" (BE). Indeed when stimulus discriminability increases, making the decision easier, the two accumulators will diverge quite rapidly and at the moment of the choice they will be far away form each other. Moreover when reaction time increases, the fluctuations of the input can push the wrong accumulator towards the threshold and therefore the two accumulators are more likely to be closer to each other. Thereby the model can account also for the inverse relationship between reaction time and confidence.

The model can work in two different decision scenarios: the optional-stopping task and the time-limited sampling task. Optional-stopping refers to the case in which subjects make the decision when they want, while in time-limited sampling the decision time is fixed by the experimenter. Using these two decision rules the DDM of Vickers (1979a) can account also for the contrasting results of Irwin et al. (1956) discussed above in sec. 1.2.3. Indeed, when the accumulation process stops before the variable has reached the bound (the time-limited case), the BE increases as a function of the stopping time because the two variables start the accumulation in the same point and theirs paths diverge

slowly. However if we consider the BE when one racer reaches the bound (the optional-stopping case) and compare trials with different first passage times, we would find that the longer the time the smaller the BE.

The one dimensional DDM cannot encode the confidence in the BE since "all that is known at decision time is that the particle has reached the bound" (as recently pointed out by Drugowitsch and Pouget (2012)). However the DDM can encode the confidence in the time taken to reach the decision, as outlined by Moreno-bote (2010). Moreno-bote (2010) demonstrate that the confidence in the race model, as the probability of reaching the correct bound, can be calculated from the state of the integrator that doesn't reach the bound and the elapsed time. This result is also valid when correlation between integrators is taken into account. Therefore, for the extreme case of perfect anti-correlation (the one-dimensional diffusion process), it can be demonstrated that the decision confidence depends only on decision time. More precisely Moreno-bote (2010) predicts that the relation between confidence level and RTs is not linear in time, as suggested by the usual fit ($a/t + b$) found by many experiments (Henmon, 1911; Volkmann, 1934; Reed, 1951; Audley, 1964), but depend on the square root of time as $a/\sqrt{t}+b$. This prediction needs a carefully designed experiment to be tested since the difference between the two types of relationships is only visible with quite long RTs.

Another solution was adopted by Kiani and Shadlen (2009) in order to account for their data of the unsure option experiment described in previous section. They propose a DDM, where the choice is determined only by the position of decision variable (i.e. the boundary has the only function of absorbing the accumulation but won't determine the reaction time). The accumulation process starts as always at zero and the duration of a trial is defined externally since they model a time-limited task. During a trial the particle diffuses, pushed towards the correct side by the drift (sensory input). A time dependent uncertainty threshold marks a region around the central point of the diffusion space. At the end of the trial if the particle is found inside of this uncertain region a "sure" target response will be triggered (meaning that the confidence in the decision is too low to commit to a choice), otherwise the sign of the decision variable will decide the target of the

saccade. This model can reliably reproduce both psychophysical and neurophysiological data but authors don't provide any novel prediction in order to test the model with new experiments, a fundamental step for the development of models.

Despite the ability of reproducing data and their simplicity DDMs lack of biological plausibility and their explanational power is sometime reduced. Therefore in the following chapters we will try to produce a reliable account of decision confidence processing based on the framework of biophysically realistic ANN.

Confidence-Based Decisions

The results presented in this chapter have been published in Insabato et al. (2010).

2.1 Introduction

In this chapter we present a model that is able to account for confidence-related decisions. Indeed as shown in fig.1.1 after a decision has been taken, the confidence associated with this decision can be encoded in some area of the brain and this information can be used in order to take subsequent decisions. This model is based on the experiemental results of Kepecs et al. (2008). They described neurons in the rat orbitofrontal cortex (OFC) as encoding decision confidence Kepecs et al. (2008) in an olfactory classification task. In the task, a mixture of odors A and B was categorized as A or B, depending on which odor was predominant. The difficulty of the task could be controlled by varying the proportion of the two odors. A second experiment was undertaken in order to examine the ability of the rats to behave in accordance with the confidence-related information. In this second experiment the rat had the possibility to abort the current trial without waiting for the reward outcome and thus to start a new trial. The results showed that a second decision (about whether to abort the trial)

could be made based on confidence in the first decision, in that the rats were more likely to abort a trial if they had made an error in the odor classification. Two different decisions should then be distinguished: a first-stage decision (about the stimulus) and a second stage decision based on the level of confidence in the first decision.

Here we propose an integrate-and-fire attractor network model, shown in Fig. 2.1, to account for decision confidence mechanisms (and described in detail in tabs.4.1,2.2 and appendix A). This is a "mechanistic" biologically realistic approach, and not a "phenomenological" approach such as an accumulator or race model that accumulates noisy evidence with a linear integrator until some threshold is reached (Vickers and Packer, 1982; Ratcliff and Rouder, 1998; Ratcliff et al., 1999; Usher and McClelland, 2001; Gold and Shadlen, 2007). We show that decision confidence is an emergent property of decision-making neural networks (for well-known implementations see Wang (2002b) and Deco and Rolls (2006)) encoded in the firing rates of the neurons (the first module in Fig. 2.1). However, this does not account for subsequent confidence-based decisions as in the second experiment of Kepecs et al. (2008). For this we propose that a second decision-making network is needed (the second module in Fig. 2.1). Thus our main proposal is for a two-layer model of confidence-related decision-making, which can account for the two decision-making processes and elucidates how confidence-related decision-making mechanisms could operate in the brain.We analyze how such a system would work, analyze its properties, and show that it accounts for the neurophysiological results described by Kepecs et al. (2008). Moreover, the model leads to predictions about new properties of these neurons that can be tested.

2.2   Results

2.2.1   The Model: Network Architecture

Our model is composed of two modules (fig.2.1). Each one is an attractor neural network implementing a decision-making process. The first module is designed to make a perceptual decision (odor classification).

We will refer to it as the decision-making (first) network. The second module, which receives inputs from the decision-making network, 'decides' whether to abort the task, estimating the level of confidence in the first perceptual decision. We will refer to it as the confidence decision-making (second) network (see Fig. 2.1).

The simple scheme of a neural network implementing a decision-making process, developed by Wang (2002b), is composed of two selective pools of excitatory neurons, one non-selective pool of excitatory neurons, and one pool of inhibitory neurons. The non-selective pool represents the background activity of the neurons not responding to the stimulus. The two selective pools represent the choices. The neurons within each excitatory selective population have strong recurrent connections, and there are weak connections between the pools. When an external input is delivered to one or both of the selective pools the activity increases, causing an enhancement of the inhibition. Since the two pools are mutually connected they cooperate increasing the activity. When the the inhibitory current is sufficiently strong a competition takes place between the selective pools. One of the two pools wins the competition and ends up with a high firing rate, while the other pool ends up with a low firing rate, indicating that a decision state is reached. The balance between competition and cooperation depends principally on the parameters of the mutual connections between the pools and the recurrent connections within a pool. The evidence for each decision is applied as an external excitatory input to each population of neurons, and biases the competition in favor of one of the two pools. The arriving random spike trains (with a Poissonian spike time distribution) together with finite-size effects produce the stochastic dynamics of the network and the probabilistic decision-making, as described in more detail elsewhere (Deco et al., 2009; Deco and Marti, 2007; Martí et al., 2008; Rolls and Deco, 2010).

The two selective pools of the decision-making network are DA and DB, which become active for decision A and B respectively. During the stimulation, pool DA (DB) receives sensory information about odor A (or B) via external input $\lambda_{A, B}$. When stimulus A (or B) is applied, pool DA (or DB) will usually win the competition and end up with high firing indicating that decision A (or B) has been reached. When a mixture is applied, the decision-making network will probabilistically
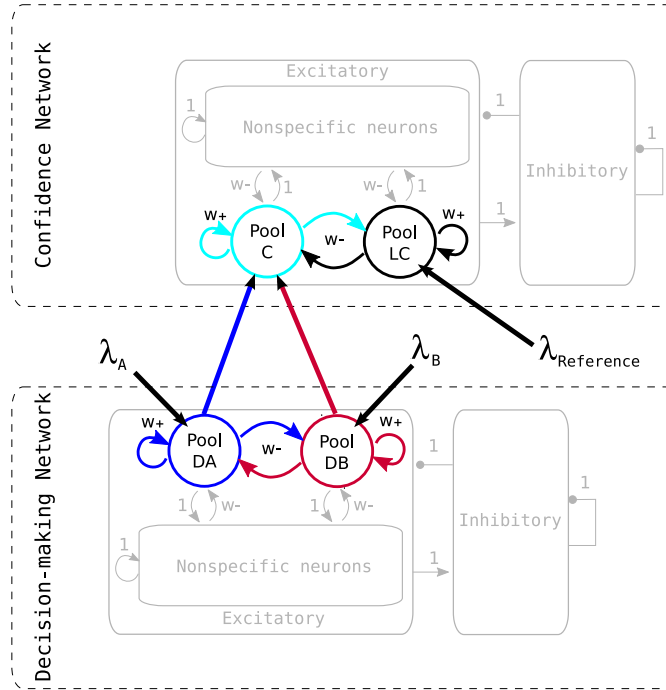
Figure 2.1: Network architecture for decisions about confidence estimates. The first network is a decision-making network, and its outputs are sent to a second network that makes decisions based on the firing rates from the first network, which reflect the decision confidence. In the first network, high firing of neuronal population (or pool) DA represents decision A, and high firing of population DB represents decision B. Pools DA and DB receive a stimulus-related input (respectively λ_A and λ_B), the evidence for each of the decisions, and these bias the attractor networks, which have internal positive feedback produced by the recurrent excitatory connections. Pools DA and DB compete through inhibitory interneurons. The second network is a confidence decision attractor network, and receives inputs from the first network. The confidence network has two selective pools of neurons, one of which (C) responds to represent confidence in the decision, and the other of which responds when there is little or a lack of confidence in the decision (LC). The C neurons receive the outputs from the selective pools of the (first) decision-making network, and the LC neurons receive $\lambda_{\text{Reference}}$ which saturates at 40 spikes/s, a rate that is close to the rates averaged across correct and error trials of the sum of the firing in the selective pools in the (first) decision-making network. In each network the excitatory pool is divided into three subpopulations: a nonspecific one, and two stimulus-selective populations. Each selective pool has strong recurrent connections (w_+), while the connections between the two selective pools are weak (w_-). All other connections are set to the default value 1.

choose DA or DB, influenced by the proportion of A and B in the mixture.

The parameters of the decision-making (first) network were chosen as follows. The inputs to the network from the sensory stimuli A and B and the synaptic weights between the neurons within a pool were set so that the network operated in a bistable regime, by which we mean that the spontaneous state is no longer stable when the decision cues are being applied. Then noise and sensory inputs bias the landscape and provoke a transition to a decision state. Thus the attractor basins are influenced by the stimulus. Whether the network can present persistent activity in the absence of stimulation, is not important for our purpose. The inputs were in addition set so that with only one stimulus in the mixture, the decision was approximately 100% correct.

The confidence network has two selective pools of neurons, one of which (C) responds to indicate confidence in the first decision and to stay with the first decision based on a level of firing from the first network which indicates high confidence, and the other of which (LC) responds when there is little or a lack of confidence in the first-stage decision. In the experiment of Kepecs et al. (2008), C corresponds to a decision to stay and wait for a reward, i.e. what they call the positive outcome population, though it really represents confidence or a prediction that the decision just made will have a positive outcome. LC corresponds to a second decision to abort a trial and not wait for a possible reward, i.e. what they call the negative outcome population, though it really represents lack of confidence that the perceptual decision just made will have a positive outcome, equivalent to confidence that the decision just made *will have* a negative outcome.

The two networks are connected by AMPA synapses that link the selective pools DA and DB of the (first) decision-making network to confidence network selective pool C. (The synaptic conductances of these connections are set to the value $g_{AMPA,ext}$=2.08 nS, but this value is not crucial for the mechanism.) The selective pool LC in the confidence network receives an external input that saturates at 40 spikes/s, in order to set the competition with pool C. This input could come from the same source at that to the C network, or could come from other brain areas, carrying for example information about the

value of different behaviors and could reflect the subject's bias in his or her confidence.

The total number of neurons in the model is $N = 2000$. For simplicity we chose to have the same number of neurons in each of the two networks $N_{\text{mod}} = N/2$. Therefore each network has $N_E = 0.8 \cdot N_{\text{mod}}$ excitatory pyramidal neurons and $N_I = 0.2 \cdot N_{\text{mod}}$ inhibitory interneurons, the proportions observed in the cerebral cortex (Abeles, 1991). The number of neurons in each selective pool is $N_A = N_B = N_C = N_{LC} = N_E \cdot f$, where f is the fraction of excitatory neurons in each selective pool. In this study we set $f = 0.15$. Each nonspecific pool in each network contains the remaining $N_E - N_A - N_B$ excitatory neurons. We modeled an equal number of neurons in each selective pool to keep the model as simple as possible, and note that equal numbers of neurons for the different attractors need not be present in this class of attractor network. Each network is fully connected, i.e. all neurons are connected to each other. We note that sparse connectivity does not change the overall dynamics of the network (i.e. the mechanism is reproducible also with a sparsely connected network), bringing about merely an increase of the noise in the network due to the finite-size effect (Mattia and Del Giudice, 2002, 2004). We make the plausible hypothesis that the connection strengths have been modified from their default value of 1 by a previous learning process, hence we set them following a Hebb-like rule, i.e. the synaptic efficacy between two cells is high if the cells had correlated activity in the past, whereas uncorrelated activity results in a weak synapse. Cells in one selective pool have strong recurrent connections w_+, while synaptic efficacy between the two selective pools is decreased, given by $w_- = (1 - fw_+)/(1 - f)$. We set these parameters (w_+, w_-) to slightly different values for the two modules. In order to achieve a better correspondence with the results of (Kepecs et al., 2008) we used a weaker w_+ in the confidence network than in the decision-making network as shown in table 4.3. During the simulation all the synaptic weights are kept fixed. All neurons receive an external input λ_{ext}, modeled as Poisson spike trains, from 800 external neurons each firing at a rate of 3 Hz, consistent with observed values in the cortex. During the stimulation the external input changes for the selective pools as described before.

Our model has a two-stage structure and assumes some type of connec-

A	Model Summary
Populations	eight
Topology	Two modules partially connected
Connectivity	full, no synaptic delay
Neuron model	Leaky Integrate-and-Fire, fixed threshold, fixed refractory time
Synapse model	Instantaneous jump and exponential decay for AMPA and GABA and exponential jump and decay for NMDA receptors
Plasticity	-
Input	Independent fixed-rate poisson spike trains to all neurons

B	Populations		
Total number of neurons	$N = 2000$	In each module	$N_{mod} = N/2$
Excitatory neurons in each module		$N_E = 0.8 \cdot N_{mod}$	
Inhibitory neurons in each module		$N_I = 0.2 \cdot N_{mod}$	
Name	**Size**	**Name**	**Size**
DA (decision A)	$N_A = f \cdot N_E$	Nonspecific (1st module)	$N_E - N_A - N_B$
DB (decision B)	$N_B = f \cdot N_E$	Inhibitory (1st module)	$0.2 \cdot N_{mod}$
C (confidence)	$N_C = f \cdot N_E$	Nonspecific (2nd module)	$N_E - N_C - N_{LC}$
LC (lack of confidence)	$N_{LC} = f \cdot N_E$	Inhibitory (2nd module)	$0.2 \cdot N_{mod}$

Table 2.1: Model summary. Network and pools details. Parameters values are given in Tab. 4.3

tivity between the two modules or stages, the decision-making network and confidence network. These modules could both be within the OFC, or they could be in different brain regions.

In the appendix B we describe the details of the mean field approximation used to determine the parameters of the synaptic strengths in the model to obtain stable operation. Other details of the architecture that was implemented are reported in Table 4.1 and shown in fig. 2.1.

C	Neuron and Synapse Model
Type	Leaky integrate-and-fire, conductance-based synapses
Subthreshold dynamics	$C_m \dot{V}(t) = -g_L(V(t) - V_L) - I_{AMPA,ext}(t) - I_{AMPA,rec}(t) - I_{NMDA}(t) - I_{GABA}(t)$
Synaptic currents	$I_{AMPA,ext}(t) = g_{AMPA,ext}(V(t) - V_E)\sum_{j=1}^{N_{ext}} s_j^{AMPA,ext}(t)$ $I_{AMPA,rec}(t) = g_{AMPA,rec}(V(t) - V_E)\sum_{j=1}^{N_E} w_j s_j^{AMPA,rec}(t)$ $I_{NMDA}(t) = \frac{g_{NMDA}(V(t)-V_E)}{1+[Mg^{2+}]exp(-0.062V(t))/3.57} \times \sum_{j=1}^{N_E} w_j s_j^{NMDA}(t)$ $I_{GABA}(t) = g_{GABA}(V(t) - V_I)\sum_{j=1}^{N_I} w_j s_j^{GABA}(t)$
Fraction of open channels	$\frac{ds_j^{AMPA,ext}(t)}{dt} = -\frac{s_j^{AMPA,ext}(t)}{\tau_{AMPA}} + \sum_k \delta(t - t_j^k)$ $\frac{ds_j^{AMPA,rec}(t)}{dt} = -\frac{s_j^{AMPA,rec}(t)}{\tau_{AMPA}} + \sum_k \delta(t - t_j^k)$ $\frac{ds_j^{NMDA}(t)}{dt} = -\frac{s_j^{NMDA}(t)}{\tau_{NMDA,decay}} + \alpha x_j(t)(1 - s_j^{NMDA}(t))$ $\frac{dx_j^{NMDA}(t)}{dt} = -\frac{x_j^{NMDA}(t)}{\tau_{NMDA,rise}} + \sum_k \delta(t - t_j^k)$ $\frac{ds_j^{GABA}(t)}{dt} = -\frac{s_j^{GABA}(t)}{\tau_{GABA}} + \sum_k \delta(t - t_j^k)$
Spiking	if $V(t) \geq V_\theta \wedge t > t^\star + \tau_{rp}$ 1. $t^\star = t$ 2. emit spike at time $t^\star$ 3. $V(t) = V_{reset}$

D	Input	
Type	**Description**	
Poisson generator	Fixed rate, N_{ext} poisson generators per neuron, each one projects to one neuron	

Table 2.2: Model summary. Neuron model and input layer description. Parameters values are given in Tab. 4.3

2.2.2 Simulation Results

Once the parameter values had been determined using the mean-field analysis (see below sec.2.4), we ran simulation trials with these values, increasing the bias value, which corresponds to altering the stimulus identity, which in this case corresponds to altering the proportion of the two odors in the mixture. On each trial the network received for the first 500 ms the external spontaneous firing level input λ_{ext}. With this input the only stable state of the network is the spontaneous activity. After that, each selective pool DA and DB received in addition to λ_{ext} the stimulus input $\lambda_{\text{A,B}} = \lambda \pm \Delta\lambda$, which provides the stimulus-specific information for the decision and drives the network dynamics. We ran different sets of simulations with $\Delta\lambda$ values of 0, 10, 20, 30 Hz. We considered that a decision was reached when the selectivity index $S = |\ln \nu_A/\nu_B|$ took a value above 1.7 and did not decrease for at least 100 ms, where ν_A and ν_B are the firing-rates of pool A and B (see sec.1.2.2 for a discussion about choice mechanisms). The same criterion was used to determine the decision in the second network using the firing rates of the C and LC pool. Pool C received the external input and the output of pools DA and DB. Pool LC received just external input for 700 ms and after that an additional input, modeled as a Poisson process of mean rate 40 Hz. We maintained the stimulation throughout the trial. This does not correspond to the experimental paradigm of Kepecs et al. (2008), but we hypothesized a working memory process upstream like the one described by Brody et al. (2003a) and Machens et al. (2005).

First we show how the firing rates of the (first) decision-making network reflect decision confidence.

Fig. 2.2c shows the proportion of correct perceptual decisions as a function of the proportion of stimulus A and stimulus B in the mixture. The decision-making is probabilistic because of the spiking-related randomness in the network (Wang, 2002b; Deco and Rolls, 2006; Rolls and Deco, 2010). Fig. 2.2a shows that on trials when the DA neuronal population which represents decision A correctly wins and has a high firing rate, the firing rate increases further with the discriminability of the stimuli $\Delta\lambda$, and thus encodes increasing confidence. The reason for the increase of firing rate with $\Delta\lambda$ on correct trials is that the exter-
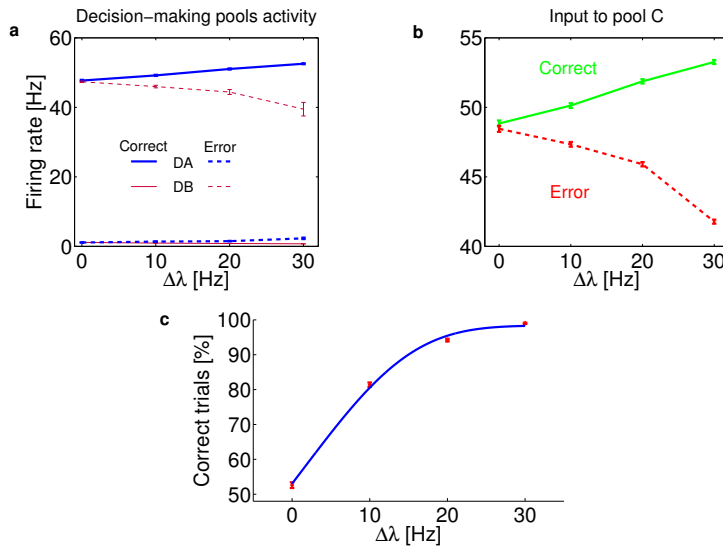
Figure 2.2: Performance of the decision-making (first) network. a) When pool DA correctly wins the competition, its firing rates are high and increase as a function of $\Delta\lambda$. When pool DB incorrectly wins making an error due to the noise and has a high firing rate, its firing rate decreases as a function of $\Delta\lambda$. This modulation of the firing rates by $\Delta\lambda$ was also observed experimentally (personal communication from Z. Mainen). The low firing rates when DB loses the competition on correct trials, and when DA loses the competition on error trials, are also shown. The error bars represent the s.e.m. (The numbers of correct trials were between 524 and 990, and the numbers of error trials from 52 to 472.) Confidence is thus encoded in the firing rates of the winning attractor, and is an emergent property of the decision-making network. The firing rates were calculated averaging over trials the activity of neurons in the last second of each trial, from time $t = 2000$ ms to $t = 3000$ ms. b) Sum of the firing rates from the DA and DB populations as a function of $\Delta\lambda$. This provides the input to the confidence (second) network selective pool C. The error bars show the s.e.m. c) The percentage correct performance of the decision-making network as a function of $\Delta\lambda$. (The error bars were estimated from the binomial distribution, and were small. The points are fitted by a Weibull function.)

nal inputs from the stimuli A or B then support the (noise-influenced) winning attractor (pool DA) and add to the firing rates being produced by the recurrent collateral connections in the winning attractor. On the other hand, on error trials the firing rates of the winning pool (now DB, which represents decision B and wins despite the evidence because of noisy firing in the network) become lower as $\Delta\lambda$ increases, because then the external sensory inputs are inconsistent with the perceptual decision that has been made, and do not support and increase the firing rate of the winning pool (Rolls and Deco, 2010; Rolls et al., 2010). (This modulation by stimulus difficulty of the firing rates of the decision-making populations was also observed experimentally by Felsen and Mainen (2009) in the superior colliculus). Confidence, which increases with $\Delta\lambda$ on correct trials and decreases with $\Delta\lambda$ on error trials (Vickers, 1979a; Vickers and Packer, 1982; Jonsson et al., 2005; Kepecs et al., 2008; Rolls et al., 2010), is thus encoded in the firing rates of the winning attractor, and is an emergent property of the decision-making network, because it was not directly implemented in the model, but arises from the simple decision process (Rolls and Deco, 2010).

Moreover, the sum of the activity of the winning and losing populations also represents decision confidence on correct and error trials, as shown in Fig. 2.2b. It is this total firing from pools DA and DB of the first, decision-making, network, which reflects decision confidence, that is provided as the input to the confidence (second) network.

We now consider the operation of the confidence decision (second) network. If the firing-rate of the winning attractor of the first, decision-making, network is high, then the confidence decision network acting as a second level network makes the second-stage decision, probabilistically as before, to have confidence in the first-stage decision, and the C population probabilistically wins the competition. If the output firing of DA and DB (reflected in their sum) is low because the perceptual decision just made has sensory inputs that are not consonant with the decision, then with weaker driving inputs to the C network, it loses the competition with LC. The confidence network in this case makes the second decision, probabilistically as before, to have a lack of confidence in the first decision, in that the LC population wins the competition. The confidence decision network thus acts as a decision-
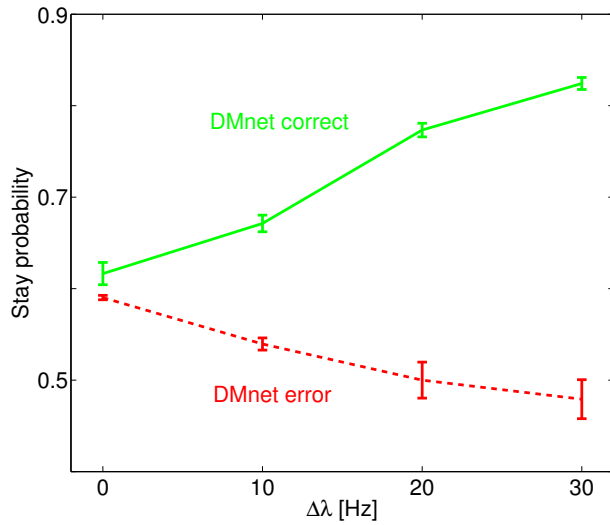
Figure 2.3: Performance of the confidence decision (second) network. The proportion of trials on which in the second network the Confidence (C) population won the competition as a function of $\Delta\lambda$ for trials on which the decision-making (first) network (DMnet) was correct or incorrect. The performance of the LC population was the complement of this. (The parameters were set so that with $\Delta\lambda$ close to 0, approximately 60% of the trials were C trials, to be qualitatively in the same direction as in the experimental findings of Kepecs et al 2008).

making network to make confident decisions if the firing rates from the first, decision-making, network are high, and to make lack of confidence decisions if the firing rates from the first, decision-making, network are low.

On trials when the (first) decision network is correct, the input to the C population coming from the DA and DB neurons increases as a function of $\Delta\lambda$, and the C pool tends to win the competition more frequently (see Fig. 2.3).

Thus a decision to act confidently about one's first decision is more likely to be made as $\Delta\lambda$ increases on correct trials. On the other

hand, when the first network makes an error, the C population tends to win the competition less frequently as $\Delta\lambda$ increases, as shown in Fig. 2.3, and correspondingly on error trials the proportion of trials on which the LC population wins increases with $\Delta\lambda$. (The percentage correct of the LC population is the complement of that shown in Fig. 2.3.) Thus a decision to lack confidence about one's first decision is more likely to be made as $\Delta\lambda$ increases on error trials, and this might make one abort such a trial, as in the experiment of Kepecs et al. (2008).

The general time structure of the neuronal activity in the model is in qualitative accordance with the experimental results (Kepecs et al., 2008). As shown in Fig. 2.4, the confidence decision takes place after the first decision, and separation of the firing rates of the two selective populations C and LC occurs after the decision-making network has reached a decision state, as in Fig. 3a-d of Kepecs et al. (2008).

It is important to examine the firing rates in the C and the LC attractor neuronal populations as a function of $\Delta\lambda$ on correct and incorrect trials, for they provide an account for neuronal responses recorded during decision-making (Kepecs et al., 2008), and those neurophysiological results in turn validate the model. We find for the confidence decision-making network that on correct trials with high $\Delta\lambda=30$ (easy perceptual decisions), C has a high firing rate, whereas it has a lower rate for $\Delta\lambda=10$, that is difficult decisions, as shown in Fig. 2.5. Conversely, on error trials when the firing rates in the first level, decision-making, network are lower, the confidence neurons C lose the competition and have relatively low firing rates, which decrease even further as the magnitude of $\Delta\lambda$ increases.

The firing rates (mean) in the confidence decision-making network of the C (confident, "positive outcome") and LC (lack of confidence, "negative outcome") populations of neurons for trials when the first decision-making network is correct or incorrect as a function of $\Delta\lambda$ are shown in Fig. 2.6.

The thick lines show the mean firing rates for the C and LC pools for all trials on which the first network was correct or in error. We identify the LC population of neurons with the negative outcome population of neurons described by Kepecs et al. (2008), which have similar prop-
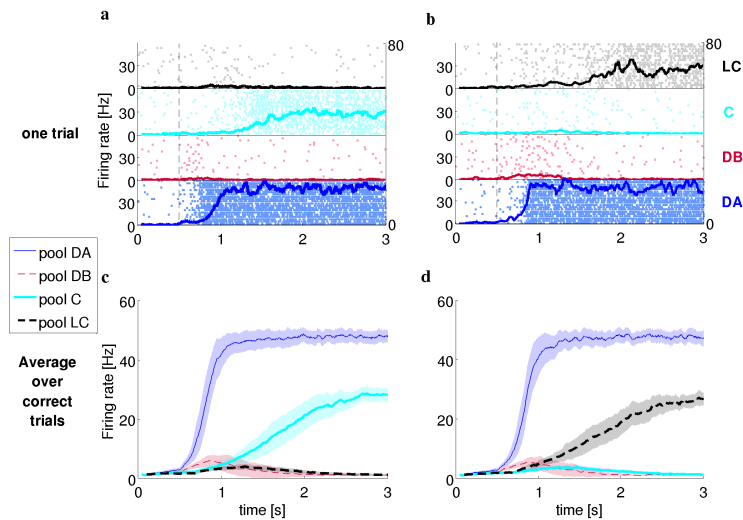
Figure 2.4: Examples of the time courses of the neuronal activity in the selective pools of the decision-making (first) network and of the confidence decision (second) network for a decision at chance ($\Delta\lambda = 0$). Panels a and b show the activity in time of selective pools in the decision-making network and in the confidence network. Rasterplots show the activity of 20 sample neurons for each selective population in one trial (dashed vertical lines mark stimulus onset). Superimposed lines show the average firing rates for that trial. On 60% of trials the confidence network selective pool C won the competition (panel a). On 40% of trials pool LC won the competition (panel b). Panels c and d show the average firing rates over all correct trials for the same conditions respectively as panels a and b. The separation of the firing rates begins after the decision is made and the general temporal structure of the network is in qualitative accordance with the experimental results of Kepecs et al (2008).
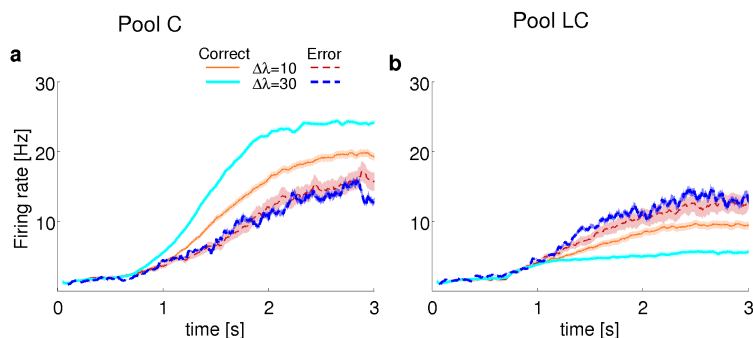
Figure 2.5: Firing rates in the confidence decision-making network of the C (confident) and LC (lack of confidence) populations of neurons for trials when the first decision-making network is correct or incorrect for easy decisions ($\Delta\lambda$=30) and difficult ($\Delta\lambda$=10) decisions. The firing rates shown are averaged merging together confident and lack of confidence trials. The same mixture is shown in thick lines in Fig. 2.6, in order to show the correspondence with the experimental results. (The shaded areas represent the s.e.m. The numbers of trials are in the range from 72 to 990). The decision cues were turned on at t=500 ms.

erties. Further, we identify the C population of neurons with the positive outcome population of neurons described by Kepecs et al. (2008), which have similar properties.

However, as shown in Fig. 2.3, the confidence decision-making network was itself increasingly incorrect (i.e. took a confidence decision that was inconsistent with the decision made by the (first) decision-making network) as $\Delta\lambda$ approached 0, and the firing rates in the thick lines of Fig. 2.6 reflect the fact that on some trials the C pool won the competition, and on some trials it lost. This is effectively how Kepecs et al. (2008) presented their data (for they knew only whether the first-stage decision itself was correct, and did not measure while recording whether the rat took a confidence-related decision to stay with or abort a trial), and there is good correspondence, as can be seen by comparing data from Kepecs et al. (2008) with Fig. 2.3. If instead of taking the mean firing rate of the C neurons based only on whether the first decision was correct, we take just the trials on which the C (confidence)
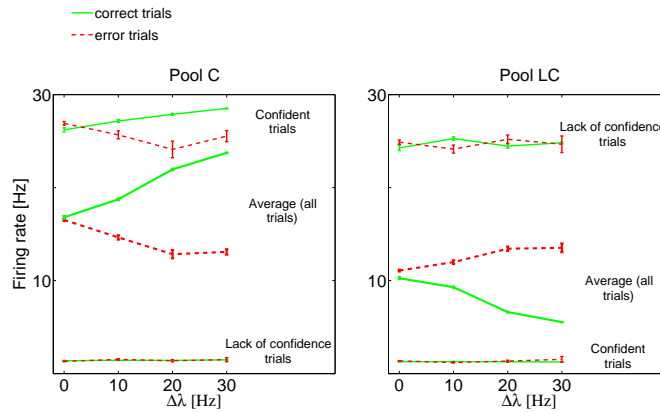
Figure 2.6: Firing rates (mean ± s.e.m.) in the confidence decision-making
network of the C (confident, 'positive outcome') and LC (lack of confidence,
'negative outcome') populations of neurons. Labels 'Confident trials' and
'Lack of confidence trials' mark respectively trials on which the C (confi-
dence) population and LC (lack of confidence) population won the competi-
tion. The activities of the C and LC populations of neurons, averaged over
all trials, confident and lack of confidence, are shown with label 'Average
(all trials)'. As shown in Fig. 2.3, the confidence decision-making network
was itself increasingly incorrect as $\Delta\lambda$ approached 0, and the firing rates
in the thick lines reflect the fact that on some trials the C population won
the competition, and on some trials it lost. (There is no correct/error dis-
tinction for trials with $\Delta\lambda = 0$. For illustration purpose, trials in which
the final choice was A were labeled as correct trials and trials in the choice
was B were labeled as error trials.) The firing rates shown were calculated
averaging the activity of neurons in the last second of each trial, from time
$t = 2000$ ms to $t = 3000$ ms. (The error bars for the thin lines (confidence
and lack of confidence decision trials) represent the s.e.m., and the number
of trials was in the range 28 to 760, with few error trials occurring with high
values of $\Delta\lambda$.)

pool won the competition, the thin lines of Fig. 2.6 show for the C pool
an average rate of close to 28 spikes/s that tends to increase with $\Delta\lambda$
when the first network is correct, and tends to decrease with $\Delta\lambda$ when
the first network is in error. (This is supported by the data shown
in Fig. 2.5.) If we take just the trials on which the C population lost
the competition, the thin lines show for the C pool an average rate
of close to 2 spikes/s. Conversely, for the LC pool of neurons, if we
take just the trials on which the LC population won the competition,
the thin lines show for the LC pool an average rate of close to 26
spikes/s. If we take just the trials on which the LC population lost the
competition, the thin lines show for this LC pool an average rate of
close to 2 spikes/s. These firing rates shown in the thin lines in Fig.
2.6 are generally as expected, and the differences with $\Delta\lambda$ are due to
whether the output of the decision-making (first) network shown in
Fig. 2.2b are consistent or inconsistent with the decision made by the
confidence decision (second) network, which is of course influenced by
the spiking noise in the confidence decision network, which can make
the wrong decision given the evidence it receives from the decision-
making network shown in Fig. 2.2b.

2.3   Discussion

In this chapter we have shown how decision confidence is an emergent
property of a neurophysiologically based decision-making process, and
is encoded in a graded way by the continuously graded firing rates of
the neurons in an integrate-and-fire attractor decision-making network.
(This is shown by the results for the first decision-making network.)
We have also shown how within this neurobiologically based framework
for decision-making, two separate networks are essential for the ability
to make a decision involving a choice about one's confidence in a prior
decision. We have also shown how the model is confirmed by and
provides a computational account for the neurophysiological findings
of Kepecs et al. (2008), and also provides a new interpretation of the
data recorded by Kepecs et al. (2008), as described in more detail next.
We also make new predictions about the types of neuronal response
that will be found when a confidence-based decision must be made, as

described next.

The fact that the changes of firing rates found in the rat by Kepecs et al. (2008) as a function of $\Delta\lambda$ are comparable with those shown in the thick lines in Fig. 2.6 provides good support for the present model. However, Kepecs et al. (2008) did not distinguish trials in which a second-layer confidence decision network was in error or not as they did not record neuronal activity when they could examine whether the rat aborted a trial, and we suggest that it would now be interesting to do this. Further, it is notable that the change of firing rate with $\Delta\lambda$ found in the rat matches only that of the thick lines in Fig. 2.6 which includes all trials irrespective of the decision made by the confidence decision (second) network, and not by the thin lines in Fig. 2.6 which reflect the decision made by the confidence network. This leads to the novel prediction that different results will be found to those presented by Kepecs et al. (2008) if in a future experiment the responses of similar neurons are separated according to whether each trial is aborted or not. We predict in particular that the neurons will have activity like that shown in the thin lines in Fig. 2.6, and will be of two types. One type will be similar to that of the C (confident in the prior decision) neurons shown in Fig. 2.6 in which the firing rate is high on trials on which the confidence decision is to stay with the first decision, and low if the confidence (second) decision is to abort the trial. The prediction further is that the firing rates of these confidence neurons will change with $\Delta\lambda$ as shown by the thin lines in Fig. 2.6a, that is these high firing rates will tend to increase as a function of $\Delta\lambda$ if the first decision (made by the decision-making, first, network) is consistent with the evidence (i.e. correct), as shown at the top of Fig. 2.6a, and to decrease as a function of $\Delta\lambda$ if the first decision (made by the decision-making, first, network) is inconsistent with the evidence (i.e. is an error), as also shown at the top of Fig. 2.3a. The second type of neuron will be similar to that of the LC (lack of confidence in the prior decision) neurons shown in Fig. 2.6b in which the firing rate is high on individual trials on which the confidence decision is to abort the trial, and low if the confidence (second) decision is to stay with the first decision. (The firing rates of the LC population do not change much with $\Delta\lambda$ as shown by the thin lines in Fig. 2.6b because the input from the saturating neurons has a fixed firing rate.) It is only

when we categorise the neurons according to whether the first decision
was correct or not that curves similar to those shown by thick lines
in Fig. 2.6 and as reported by Kepecs et al. (2008) will be found, and
such curves and analyses do not capture fully the properties of the
confidence decision-related neurons, which are as shown in the thin
lines in Fig. 2.6.

We used the model to account for neural data recorded in the OFC
by Kepecs et al. (2008). However, different brain areas are involved in
different types of decision-making (Kim and Shadlen, 1999; Hernandez
et al., 2002; Romo et al., 2002a; Rolls and Deco, 2010), and it is accord-
ingly plausible that other brain areas can process confidence-related in-
formation. Therefore we tried to keep the model as simple as possible
to propose a generic mechanism for confidence-related representation
and decision-making, that is not only consistent with and provides
an interpretation of experimental data from Kepecs et al. (2008), but
also provides a generic account of confidence-related decision-making
in other brain areas.

Our model is designed based on a binary decision-making task, but it
could be slightly modified to encompass also multiple choice decision-
making. In some recent work the theoretical framework of biased com-
petition that we adopted has been developed to account for multi-
ple choice decision-making (Furman and Wang, 2008; Albantakis and
Deco, 2009b). A possible extension of our model would be a com-
bination of an architecture like the one proposed by Albantakis and
Deco (2009b), but with a second confidence network strongly connected
with the decision neurons. We propose that such an extended model
based on our idea can also account for decisions based on confidence
in multiple-choice decision-making processes.

The architecture of our model is based on two layers, one for the per-
ceptual decision, and one to monitor the confidence level of the first
decision. However there is no restriction to two layers, and deeper ar-
chitectures could be built to perform more complex functions. In fact,
the second layer also undertakes a decision process, and hence a third
layer could monitor the activity of the second layer. The proposed
mechanism could thus be extended to account for nested monitoring
functions (cf. Hofstadter (2007)). Although no natural restriction is

imposed on the mechanism, and learning could shape a nested network hierarchy, a second-order confidence-related decision (and eventually a third-order) may be of less use than a judgement about a first-level (e.g. perceptual) decision.

The confidence decision (second) network is in effect monitoring the decisions made by the first network, and can cause a change of behavior, choosing to abort the trial, if the second network's assessment of the decision made by the first network is that the first decision is not a confident decision. Now this is the type of description, and language used, to describe 'monitoring' functions, taken to be high level cognitive processes, possibly related to consciousness Lycan (1997); Block (1995). For example, in an experiment performed by Hampton (2001a) (experiment 3), a monkey performing a short-term memory task could choose an 'escape flag' to start another trial. With longer delays, when memory strength might be lower partly due to noise in the system, and confidence therefore might be lower, the monkey was more likely to choose the escape flag. The experiment is described as showing that the monkey is thinking about his own memory, that is, is a case of meta-memory, which may be related to consciousness (Heyes, 2008). However, the decision about whether to escape from a trial can be made just by adding a second decision network to the first decision network. Thus we can account for what seem like complex cognitive phenomena with a simple system of two attractor decision-making networks (Fig. 2.1). The design of Kepecs et al. (2008) was analogous, in that the rat could choose to abort a trial if decision confidence was low, and again this functionality can be implemented by two attractor decision-making networks, as described here.

There are other more complex types of 'self-monitoring', such as is described as occurring in a commentary that might be based on reflection on previous events, and appears to be more closely related to consciousness (Weiskrantz, 1997; Rolls, 2007). Our aim was not to account for such complex monitoring functions. Rather we claim that some types of 'self-monitoring' are computationally simple, and the model we propose can be a building-block for a better understanding of such a high level cognitive function.

As a last remark we want to stress our belief that models and theories

have to be judged and selected on the basis of their capability to predict new results. These predictions can then be tested in new experiments and thereby the models can be validated or falsified. It is with this spirit that we undertake in the next chapter an analysis of neural data recorded in behaviong monkeys during a decision-making task. In the next chapter indeed we will present novel neurophysiological results that, on one hand, confirm the model proposed in this chapter and, on the other, adds new evidences on top of our understanding of the neural substrate of decision confidence.

2.4   Methods

2.4.1   Model Details and Mean-Field Reduction

Detailed description of the neuron and synapse model is given in the appendix A and in tabs. 4.1 and 2.2.

We use a mean-field reduction (Brunel and Wang, 2001b) in order to study the space of the principal parameters of each network.

The mean-field approximation allows reduction of the number of dynamical variables, by describing the average firing rate of each neuronal pool in the limit of an infinitely large number of neurons. The network dynamics could converge to one of four attractors: a spontaneous state, where the selective pools have low activity; two selective states with one selective pool firing at a high level and the other inhibited; and a mixed state with both pools highly active. With some parameter values the network could sustain just one stable state, or could show bistable, behaviour, or could show multistable behaviour in which the spontaneous state and each of the two decision states are all possible stable states when the inputs to the network are being applied (Brunel and Wang, 2001b). Details of the mean-field analysis are given in the appendix B.

We extensively explored the parameter space for λ and $\Delta\lambda$, using few different values of w_+. All other parameters were set accordingly to the results of ?. Our goal was to find the parameters that provide regions of bistability (in which each of the decision states is stable) or
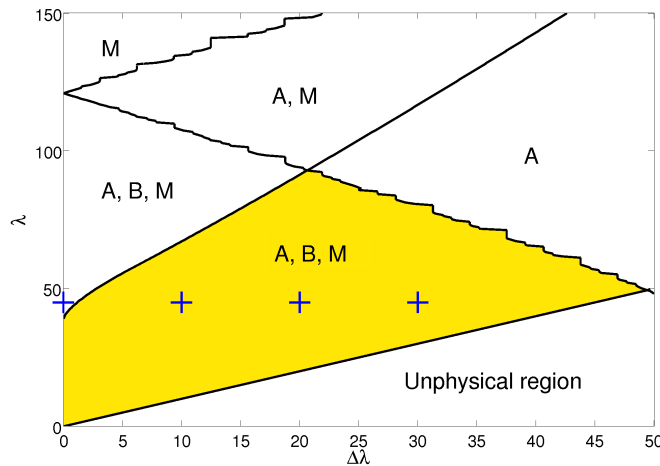
Figure 2.7: Mean field bifurcation diagram. The different areas represent different region of operation of the decision-making network. The letters indicate network stable states: A, population DA wins the competition; B, population DB wins the competition; M, population DA and DB end up with high firing rates indicating that no decision has been reached (these trials were not included in the analysis). We were interested in the multistability region and, in particular, in the bistable region (the shaded yellow area). The model is robust to a wide parameter space (indicated in yellow), and other parameters in the multistable region indicated by labels A, B, and M can be used. The crosses indicate the chosen working points where full spiking activity was investigated.

multistability (in which each of the decision states and the spontaneous state is stable) for the decision-making network. As shown in fig. 2.7 there is a wide region of parameter space where the network operates in a bistability regime. This shows a good robustness of our results with respect to these parameter variations. We chose our working point at λ=45 Hz. Once the parameters of the two separate modules had been set, we chose the values of the intermodule connections to be the same as those of the standard connections from the external inputs. All parameter values are reported in Table 4.3.

Parameter	Value	Parameter	Value
C_m (excitatory)	0.5 nF	V_E	0 mV
C_m (inhibitory)	0.2 nF	V_I	-70 mV
f	0.15	V_L	-70 mV
$g_{AMPA,ext}$ (excitatory)	2.08 nS	V_{reset}	-55 mV
$g_{AMPA,ext}$ (inhibitory)	1.62 nS	V_θ	-50 mV
$g_{AMPA,rec}$ (excitatory)	0.104 nS	w_+(DM net.)	1.8
$g_{AMPA,rec}$ (inhibitory)	0.081 nS	w_+(Conf. net.)	1.7
g_{GABA} (excitatory)	1.287 nS	α	0.5 ms^{-1}
g_{GABA} (inhibitory)	1.002 nS	$\lambda_{\text{Reference}}$	40 Hz
g_{NMDA} (excitatory)	0.327 nS	λ_{ext}	2.4 kHz
λ	45 Hz	$\Delta\lambda$	[0 30] Hz
g_{NMDA} (inhibitory)	0.258 nS	τ_{AMPA}	2 ms
N_E	800	τ_{GABA}	10 ms
N_I	200	$\tau_{NMDA,decay}$	100 ms
N_{ext}	800	$\tau_{NMDA,rise}$	2 ms

Table 2.3: Default parameters used in the simulations.

2.4.2 Implementation

Once the parameters were fixed using the mean-field analysis we ran spiking simulations. Both the mean-field reduction and spiking simulations were implemented in custom C++ programs. For the mean-field numerical integration we used an Euler routine with a step size of 0.1. For the spiking simulations we used a second-order Runge-Kutta routine with a time step of 0.02 ms to perform numerical integration of the coupled differential equations that describe the dynamics of all cells and synapses. The population firing rates were calculated by performing a spike count over a 50 ms window moved with a time step of 5 ms. This sum was then divided by the number of neurons in the population and by the window size.

Confidence neurons in the primate brain

The results of this chapter are included in Martinez-Garcia et al., Neural correlates of decision confidence in monkey prefrontal cortex. Submitted.

3.1 Introduction

In this chapter we present neurophysiological data underlying decision confidence processes. In chap.1 we reviewed the existing neurophysiological results about decision confidence. However all these results leave still as an open question how the neural signal of confidence is encoded in single trials.

In particular, it remains unclear whether neurons encode confidence in a continuous manner, or in a discrete manner. The results of Kepecs et al. (2008) seem to suggest that OFC encodes confidence in a continuous way, which has also been suggested by theoretical studies based on diffusion-like models (DDM) (Vickers, 1979b; Moreno-bote, 2010; Drugowitsch et al., 2012). These theoretical studies seem to indicate

that confidence is encoded by the position of accumulators at the moment of choice (race models), or otherwise by the time it takes the system to reach a decision (diffusion models). A binary coding of decision confidence, on the other hand, is suggested by the activity of LIP; Here neurons present a very different firing-rate in confident trials, as compared to uncertain trials (Kiani and Shadlen, 2009). Furthermore, it is interesting to note that Graziano and Sigman (2009) have shown that human subjects' confidence ratings presented a highly bimodal distribution. A recent biologically realistic attractor model (Insabato et al., 2010) combines continuous and discrete encoding into a two-stage model. The authors suggest that continuous confidence signals encoded in the decision-making neurons are translated into a binary response of confidence neurons, which discriminates high from low confidence. A similar mechanism has also been proposed within the DDM-framework of (Pleskac and Busemeyer, 2010).

In the present chapter we have looked at how confidence is represented in single neurons recorded from ventral premotor cortex (PMv) while monkeys perform a visual discrimination task (Pardo-Vázquez et al., 2008). Our aim is to shed light on the encoding mechanisms of decision confidence in the primate brain. PMv neurons seem well suited to evaluate decision confidence, since previous studies have shown a central role of premotor cortex in the conversion of a decision into an action (de Lafuente and Romo, 2006; Hernández et al., 2010). More precisely, it has been found that PMv neurons encode higher cognitive processes, such as decision-making (Romo et al., 2004; Pardo-Vázquez et al., 2008), and performance-monitoring (Pardo-Vázquez et al., 2008, 2009). We have analyzed neuronal activity both across trials using Linear regression (e.g. Hernández et al. (2010)), and in single trials, using the Hidden Markov model (HMM). As shown in Pardo-Vázquez et al. (2008), in PMv decision-making neurons exist that, besides the response, also encode the difficulty of the decision. The pattern displayed by these neurons is similar to that of decision-making neurons in the model of Insabato et al. (2010). Therefore, the dataset as recorded by Pardo-Vázquez et al. (2008) is particularly well-suited to look for confidence encoding neurons. Indeed our results show for the first time in the primate brain a pool of neurons whose mean firing rate over trials continuously encode confidence; a result similar to what

presented by Kepecs et al. (2008) in rats OFC. Interestingly we also found neurons showing two distinct levels of activity over trials. Thus our results suggest that both continuous and discrete coding schemes for confidence are active in the brain.

3.2 Results

3.2.1 PMv Neurons Encode Decision Confidence

We studied the decision-process in the primate brain during a simple binary decision task. Two male monkeys (Macaca mulatta) performed a two-interval two-alternative discrimination task. They had to compare the orientation of a reference bar, presented during the first interval, with that of a test bar, presented during the second interval. They then had to decide whether the test bar was tilted right or left as compared to the reference bar (see Fig. 3.1 and Methods for details). The level of difficulty of the task was controlled by varying the difference between the orientation of the first and the second bar, i.e the test bar's relative orientation (TRO). The TRO was varied from one up to four degrees and in both directions.

Single cells from PMv were recorded while monkeys performed the task. For a more detailed description of the task, behavioral results and neural recordings see Methods and Pardo-Vázquez et al. (2008).

While Pardo-Vázquez et al. (2008) describe neurons in PMv that encode subjects choice, our principal objective was to find neural signatures of decision confidence computations in this area of the primate brain. It is plausible that confidence-related computations take place in the same area as where the decision is encoded, given the dependence of decision confidence processes on decision-making processes. In addition, Pardo-Vázquez et al. (2008) also found decision-making neurons that encode the difficulty of the decision, a computation that is fundamental to confidence processing. And Kiani and Shadlen (2009) found a correlate of decision confidence in the same neurons that encode the choice in monkeys' LIP. We therefore analyzed the activity of PMv activity recorded during the decision task above described.
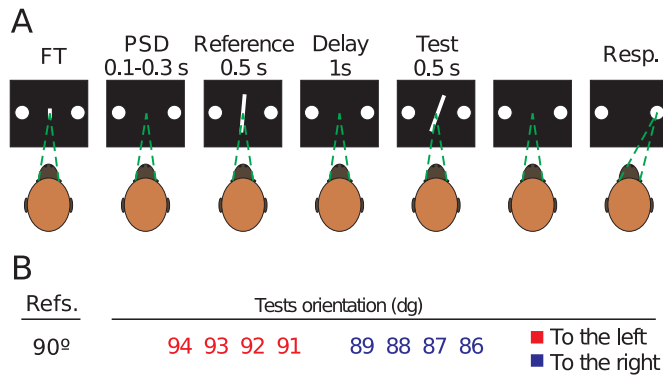
Figure 3.1: Experimental paradigm. A: The trial starts when the monkey acquire fixation to the central target (FT). A brief pre-stimulus delay follows (PSD). The reference bar is presented for 500 ms with one of three possible orientations (Reference). During the subsequent delay the subject has to mantain fixation (delay). The test bar is shown tilted to the left or to the right respect to the reference bar. The orientation of the test bar relative to the reference manipulates the difficulty of the trial. When the test bar disappear the subject can decide whether the test was right or left tilted respect to the reference by making a saccade towards the rigth or left choice targets respectively. B: Distribution of the possible orientations of the test bar. For more details see Pardo-Vázquez et al. (2008).

Our analysis was restricted to a subset of the recorded neurons (336 neurons, see Pardo-Vázquez et al. (2008)), comprising the cells that were relevant to the decision task. Unless specified otherwise, in the following we will only describe correct trials, since error trials were enough only for few neurons. We identified a population of neurons (49 cells) whose firing-rate was high for both right and left decisions (as can be seen in the raster plot of a single neuron in Fig. 3.2A). Even if not predictive of the choice, the firing-rate of these neurons encoded the difficulty of the task, independently of the subjects choice, as revealed by the linear model (LM) we used (see Methods).

Fig. 3.2B shows the evolution in time of the coefficient of the LM for a single neuron. The shaded area marks the time-window where the coefficient of the LM was significant, giving an estimate of the encoding time-window of the neuron. For the neuron in Fig. 3.2 the encoding
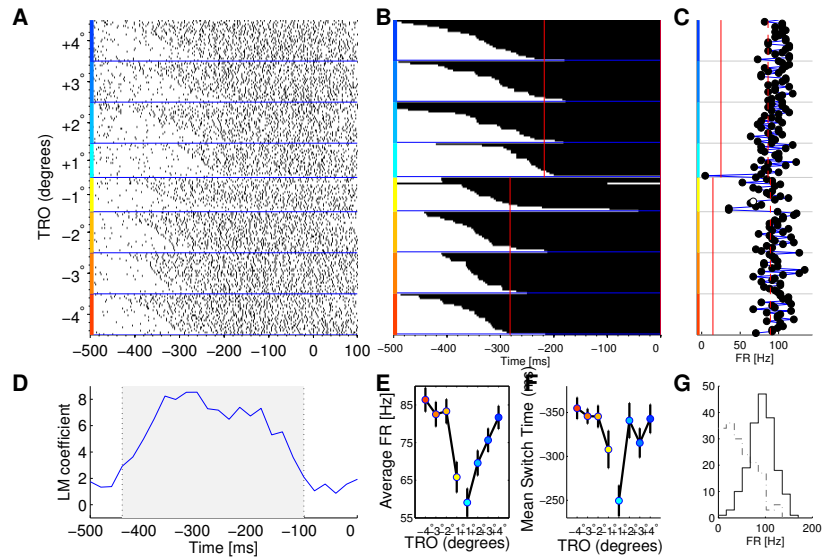
Figure 3.2: Single neuron from PMv cortex enconding confidence in a continuous way. Each different test relative orientation (TRO) is assigned to a different color: colors form dark blue (red) to ligth blue (yellow) correspond to right (left) responses with increasing difficulty, i.e. relative orientation from 4^o (-4^o) to 1^o (-1^o). Trials are aligned to the saccade; around -510 ms the second bar was shown. A: Raster plot. The trials are sorted by the TRO and according to the timing of state switch (as indicated by the HMM). B: Time course of state switchings according to HMM for the same neuron. Every row represents a trial. State one is represented by color white, while black represents state two. Trials start in state one and later change to state two (indicating that they increased their firing rate). The vertical red line indicates when the 90% of the trials changed state. C: Time averaged firing rate of the singles trials taken in the window from the red line of panel B to the saccade; red lines represent the mean firing rate over trials for "down" state (solid) and "up" state (dahsed). D: The time evolution of the LM coefficient. Shaded area corresponds to the period in which the coefficient was significative ($p < 0.05, tstatics$). E: Average firing rate as a function of TRO. The time average is taken during the time window marked by the shaded area D. Errorbars respresent SEM. F: Mean switch time (according to HMM analysis) as a function of TRO. G: Histograms of the firing rates in the "down" (dashed) and "up" state (solid).
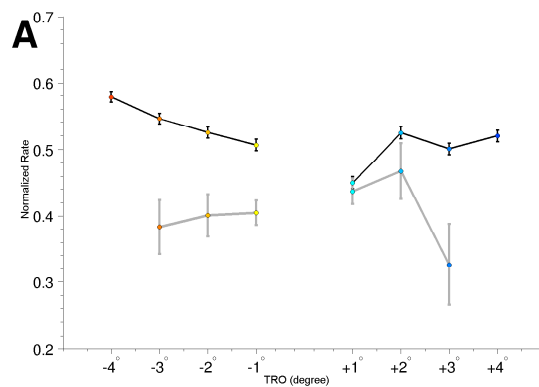
Figure 3.3: Average population activity. Normalized firing rate as a function of TRO of the confidence neurons that present a positive relation between firing rate and difficulty. Black lines: correc trials; grey lines error trials. Each neuron firing rate was normalized to its own maximum firing rate. Bars represent SEM.

window spanned approximately 300 ms. As shown in Fig. 3.2C, during this period the firing-rate of the neuron, as a function of TRO, increases for both positive and negative values of TRO, producing a v-shaped pattern. This pattern can also be seen in the population activity (Fig. 3.3, black line).

Moreover, by pooling data from all the neurons we were able to analyze firing-rates during error trials. When the behavioral response was incorrect the firing-rate of this population showed an inverse pattern compared to correct trials. Overall, the normalized firing-rate separated in correct and error trials formed an x-shaped pattern, which had already been described by Kepecs et al. (2008) in rats' OFCs as a correlate of decision confidence. As has already been remarked, consistent neural recordings during error trials were rare, but in order to confirm the pattern of the pooled responses at the level of single neuron we analyzed fifteen of the 49 neurons for which we had enough error trials recordinged (see Methods for details). To check whether the x-pattern was present in the single-cell activity we looked for neurons having the LM coefficients significant, but with a different sign in error trials compared to correct ones. According to this analysis the

coefficient was statistically significant for five of the fifteen neurons.

The x-pattern associated with confidence is poorly understood in terms of the computations from which it could be said to arise. In the following, we will try to shed light on this matter.However, we only analyze correct trials because of the small number of errors in most of the recordings. The increased firing-rate as a function of the absolute value of TRO, i.e., the v-shaped pattern associated with correct trials, can arise from at least three distinct mechanisms (for a pictorial representation see Fig. 3.4). 1) Rate coding (panel A): neurons increase the firing-rate respect to the baseline in proportion to the confidence in the decision (something which highly correlates with the difficulty of the trial (Vickers, 1979b). 2) Switch time coding (panel B): neurons increase the firing-rate, switching from a low to a high activity state, with a different timing according to the confidence, and with the average rate reflecting this timing. 3) Binary coding (panel C): neurons have a binary response, i.e., they increase the firing-rate only in high confidence trials (whilst when confidence is low they remain in a down state). In this last scenario the proportion of confident trials depends on its level of difficulty, and mixing trials of high and low activity produces the v-shaped pattern of average firing-rates (Insabato et al., 2010).

In order to identify neurons implementing each of these mechanisms we used different statistical techniques. Although we present them here as separated mechanisms, we do not rule out the possibility that they could all appear at the same time.

We started by verifying whether the switch timing had any relevant effect in our data. To do so we used a Hidden Markov Model (HMM) (for its application with single neuron recordings see Ponce Alvarez et al. (2008)) which is able to detect when a system switches from one state of activity to another (see Methods for details). In Fig. 3.2B we show a summary of the two-state HMM analysis for one confidence neuron (each row represents a trial). The color of the row changes from white to black when the neuron goes from a low to a high-activity state. This neuron exhibits a lot of variability in the switch timing, changing state from just a few milliseconds up to 300 ms after stimulus onset. The timing of the change was correlated with the difficulty of
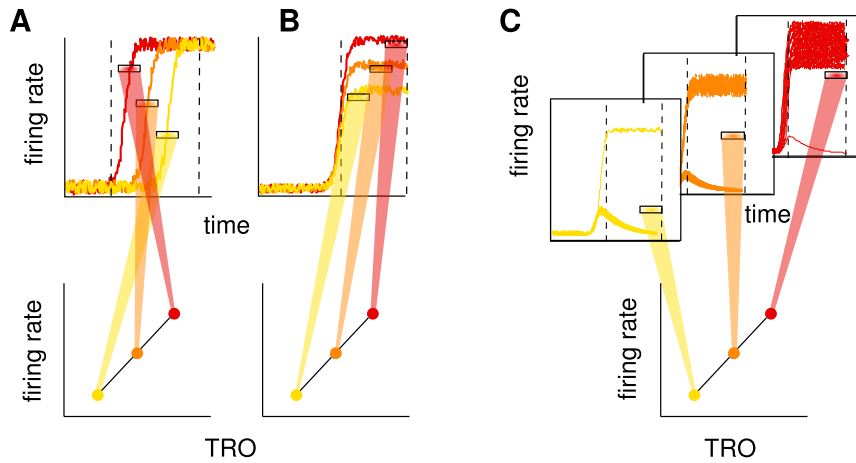
Figure 3.4: Pictorial representation of possible mechanisms underlying the confidence "X" shaped pattern. The figure only illustrate a linear relation between firing rate and test relative orientation (TRO) since it is the basis of the "X" pattern. A: switch timing code; in the upper panel the time evolution of firing rate is shown for three trials (one for each TRO; the color code is the same of the bottom panel). Each trial present a switch from a low activity state to a high activity state. The three horizontal marks show the time averaged firing rates taken in the window enclosed in the vertical dashed lines. In the bottom panel this average firing rates are shown as a function of TRO. The different switching time of the trials produces different firing rates. B: rate code; each trial reaches a different level of firing rate in the high activity state (upper panel) and this is reflected in the mean firing rate (bottom panel). C: binary code; only some trials switch to the high activity state while others remain in the "down" state. The number of trials that switch state depends on TRO. When many trials are in a "up" state (red trials) the mean over trials of the time averaged firing rate is higher respect to the case of many trials in the "down" state (yellow trials).

the trial (Kendalls correlation coefficient $\tau = 0.18, p < 0.05$). Fig. 3.2F represents the mean switch time as a function of TRO. Once we had determined when a neuron changes its state we were then able to assess the relevance of the rate coding mechanism. The firing-rate after the state switch is represented in Fig. 3.2C for the same confidence neuron. Each dot represents the time-averaged firing-rate of one trial, color-coded according to the state assigned by HMM (for comparison the dashed and solid red lines represent the firing-rate of the high and low states founded by HMM respectively).

To estimate whether the increase in firing-rates was proportional to the difficulty (i.e., the rate coding mechanism of Fig. 3.4B), we first calculated the average firing-rate from when 90% of the trials switched states (red vertical line in Fig. 3.2B), until the coefficient of LM had a significant value (shaded region in Fig. 3.2D). Then we effectuated a correlation analysis between the level of difficulty and the average firing-rate. We obtained a significant correlation coefficient for the neuron in Fig. 3.2 ($\tau = 0.24$, Kendalls correlation, $p < 0.05$), which suggests that it could be the firing-rate of the neuron in the up state that encodes the trial's level of difficulty.

To summarize, we found that twelve neurons presented a significant impact on the timing in the formation of the pattern, while nine neurons increased the firing-rate proportionally to the difficulty of the trial, thereby implementing the rate coding mechanism. There were also five neurons that presented both switch timing code and rate code (see Fig. 3.6 for a graphical representation of all classes of neurons). We note that we could apply this method only to 28 out of 49 confidence neurons, as we considered the HMM analysis was only reliable under some constraints (see Methods).

3.2.2 Discrete Confidence Encoding

The binary mechanism postulated above corresponds to a discrete confidence encoding like the one hypothesized in Insabato et al. (2010). Although a continuous representation of confidence is probable at some stage of the sensory-motor integration, a discretization stage is needed to account for the behavioral effect of the confidence computation

(Pleskac and Busemeyer, 2010). Indeed, both the usual confidence ratings and confidence-related decisions (Kepecs et al., 2008) require the selection of different alternatives. In order to identify neurons with a binary response we hypothesized that the distribution over trials of the mean firing-rate as calculated during the test-bar presentation, has to consist of two different distributions. The resulting distribution is not necessarily bimodal but it should differ substantially from the expected Poisson distribution (Softky and Koch, 1993; Bair et al., 1994; Shadlen and Newsome, 1994). For each trial, therefore, we took the average firing-rate over a 200 ms time-window, ending at the time of subject's response. Then we fitted these mean firing-rates to the average of two gamma distributions, parametrically varying the shape and the mean of the distributions:

$$B = (F_\Gamma(x; k_1, \mu_1) + F_\Gamma(x; k_2, \mu_2))/2,$$

where $k_{1,2}$ is the shape parameter and $\mu_{1,2}$ the mean. We used a gamma distribution because of its broad generality. We used a chi-squared test (χ^2) to evaluate the goodness of fit for each set of parameters (see Methods for details). Fig. 3.5F shows the histogram of the firing-rate of one single neuron (blue line) and the best fit model (black line), composed of the two gamma distributions (red line).

In our analysis we took all the possible combinations of the parameters of the two distributions into consideration. Therefore, it is possible that, if the actual empirical firing-rate distribution is unimodal (e.g. Poisson, normal, etc.), it could be well fitted by the mean of two gamma distributions with similar parameters. In order to eliminate this possibility we tested whether a model with four parameters (two for each distribution) was more adequate than a model with only two parameters (only one gamma distribution). To this aim we used the Bayesian Information Criterion (BIC) that, while comparing the likelihood function of the two models, corrects the result by penalizing for the number of parameters. Therefore, even if the likelihood of the single distribution model were equal to that of the double distribution model, the BIC would always prefer the simpler model (or, conversely, a double distribution model would be preferable only if it was able to explain much more than the single distribution model).
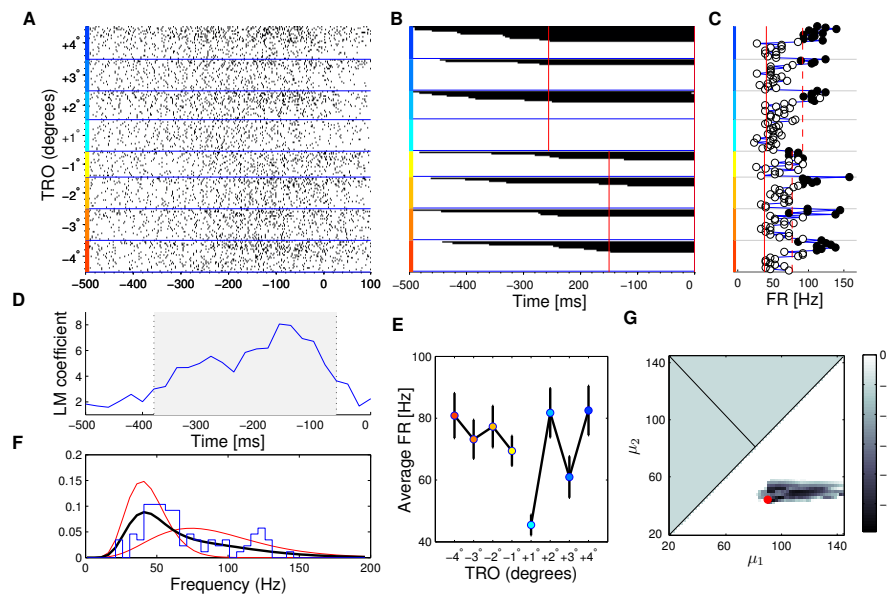
Figure 3.5: Single neuron from PMv, implementing a binary confidence encoding. Color and line convention is the same as figure 3.2. A: Raster plot. The trials are sorted as in 3.2. B: Time course of states switching, according to HMM. Same conventions as in fig.3.2B. In contrast to neuron depicted in figure 3.2 not every trial of this neuron change from state one to state two. This clearly indicates that the neuron has two distinct behaviors: in some trials it increases the firing rate while in others it remains at a lower level of activity. C: Average firing rates of single trials. D: LM coefficient value. Shaded area corresponds to the period in which the coefficient was significative ($p < 0.05 tstatics$). E: Average firing rate versus TRO. The time average was calculated during the period marked by shaded area in D. F: Histogram of firing rate during the 200 ms previous to the saccade (in blue). In black the distribution that best fits the data. This distribution is the average of two gamma distributions shown in red. The fitting has been computed separately for positive and negative TRO, the picture corresponds to negative TRO values. G: χ^2 goodness-of-fit. The fitting was done in the four dimensions of model parameters but here we show the results in the plane of μ_1, μ_2. Each point in the plane correspond to a couple of μ_1, μ_2 values. The color of the point represents difference between the highest acceptable p-value (0.05) of the χ^2 and the p-value obtained with that set of parameters values. See methods for details. The white area indicates regions of the parameters values that give non significant results. The red point represents the best fit.
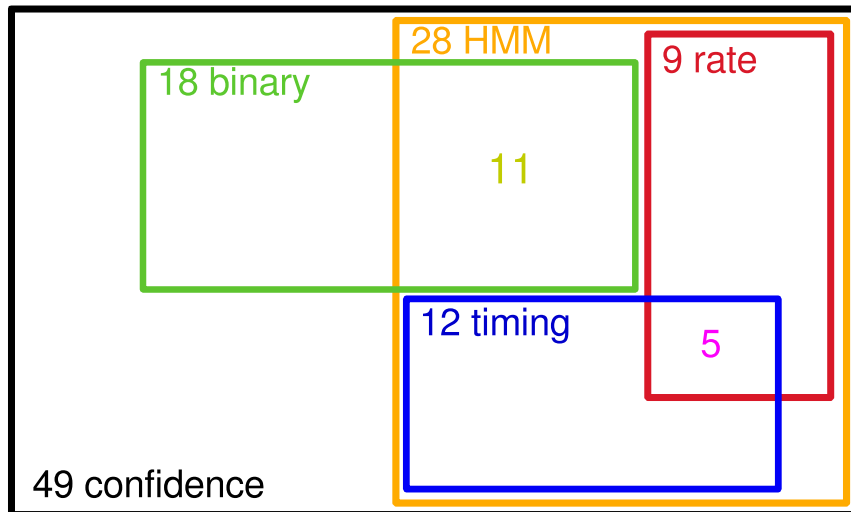
Figure 3.6: Graphical representation of the different classes of neurons. The label and number in each rectangle indicate the class and the number of neurons in that set. Out of the 49 confidence neurons only 28 could be alayzed with HMM, as explained in the text.

In conclusion, we consider a neuron to have given a binary response if the chi-squared test gives a significant result and the model with a double distribution is better than the one with only a single distribution according to the BIC. Fig. 3.5G represents the goodness of fit (for significant values only) in the space of the means of the two gamma distributions (μ_1, μ_2). In this space the color of each point represents the difference between the highest acceptable value of the probability of the chi-square statistics $(p < 0.05)$ and the actual value of this probability obtained for the combination of μ_1 and μ_2 of that point. It is interesting to note that the points where we get significant values do not lay on the diagonal (where the two means are identical). We ran this analysis on the entire confidence related population (49 neurons) and we found that eighteen neurons displayed a binary encoding of decision confidence in the case of at least one behavioral response (e.g., left) (see Fig. 3.6 for a graphical representation of all classes of neurons).

In these neurons the v-shaped pattern of the firing-rate is the result of the fact that the proportion of trials with high firing-rate correlates

with the difficulty of the trial. We also reasoned that the neurons showing a binary behavior should also lead to a characteristic pattern showing up in the HMM analysis: they should present a state switch only on a subset of trials. And indeed this pattern can be seen in Fig. 3.5B. In Fig. 3.5C the mean firing-rate of the neurons is shown for each trial, separated in high firing-rate trials (filled dots) and low firing-rate trials (open dots). The separation of firing-rates can be clearly seen when compared to the continuous confidence encoding neurons (see Fig. 3.2C).

3.3 Discussion

In this chapter and we address the question of how neurons in the primate brain encode decision confidence in single trials. While evidence found in rat OFCs seems to suggest that confidence is encoded by the continuously varying firing-rate of neurons, recordings made in monkey LIPs show that confident trials and uncertain trials do not display the same pattern. Continuous encoding schemes and discrete encoding schemes involve different computations and probably serve different functions; it is therefore important to know which of the two (or both) is implemented in the brain.

We have demonstrated that during correct trials neurons in primate PMv increase their firing-rate as a function of stimulus discriminability (in our experiments: the relative orientation of a bar), whereas in error trials the firing-rate decreased. This peculiar pattern has already been described by Kepecs et al. (2008) as a correlate of confidence in the rat OFCs. Here, for the first time, a similar result in monkey PMv cortex is presented.

It is worth noting that the pattern emerges when the firing-rates of neurons over several trials, and with the same discriminability, are averaged together. Nonetheless, different computations performed by neurons in single trials can produce the same pattern of average firing-rates. We suggested three hypothesis: 1) The switch time coding: when the activity of the neuron changes, the difficulty of the decision, is encoded in the timing of the change, 2) The rate coding: the dif-

ficulty is encoded in the firing-rate, after the change has taken place; or 3) The binary coding: the neuron only changes activity in high (or low) confidence trials and the proportion of high confidence trials changes according to the discriminability of the stimulus. The first two alternatives correspond to a continuous encoding of confidence, whereas the last one is a form of discrete encoding. We found that, in fact, all three mechanisms are at work in monkey PMv neurons. For certain neurons the timing and firing-rate mechanisms work together, i.e., a neuron that changed state earlier on less difficult trials will also have a higher firing-rate after the change. Other neurons present a binary response (increasing activity only in some trials), which suggests a possible role in confidence judgments.

An important question is: why should neurons use different schemes to encode confidence? Our hypothesis is that confidence neurons carry-out more than one function in the sensory-motor path. It is logical that a "confidence" module would encode decision confidence on a continuous scale, since confidence is a graded sensation. However, if decision confidence is to have behavioral relevance, the information about confidence needs to be discretized (see Fig. 1.1). Our hypothesis is that, while certain neurons encode confidence in a continuous manner, other neurons read-out this scale and transform it into a discrete quantity in order to produce consistent behavior.

This idea has been partially implemented in the biologically realistic attractor neural network of Insabato et al. (2010) and in the DDM framework of Pleskac and Busemeyer (2010). Indeed, Insabato et al. (2010) have shown that the sum of the firing-rate of the decision neurons is a good representation of the confidence on a continuous scale. This representation is given as input to another decision-making network that then makes a choice based on the confidence estimation. Therefore, our results of binary confidence neurons confirm the predictions of the model about a confidence-related decision-making module. However, we must note that the model does not explicitly consider the continuous confidence representation as found in PMv neurons. In the model, confidence is represented implicitly by the sum of decision-neurons' firing-rates, as no neuron seems to encode it directly.

Despite the results as presented in this article, it could be objected

that, like in Kepecs et al. (2008), the task has no behavioral correlate of confidence, and that reasonable doubts could thus be cast on the interpretation of the results. Although we think that a set-up where both choice and confidence are recorded would better serve the scope, we would like to note that, rather more intriguingly, we found confidence encoding neurons in a context where confidence estimation was not relevant. Indeed, it stands to reason that a mechanism that estimates confidence and that makes decisions based on this would stay in place even when not in use. We therefore expect that a more suited task will lead to the recording of even stronger signals of confidence.

In that respect, we note that we always applied conservative methods in order to filter-out spurious results, probably paying the penalty of type II errors. For instance, when using the linear model on our firing-rate time series, we excluded the possibility of getting significant results in consecutive bins, by chance alone. In order to do so we used a minimum of four consecutive bins (see Methods for details). Also, when we looked for binary confidence neurons (fitted by a model with two distributions) we wanted to rule out the possibility of a simpler model (with a unique distribution) explaining the data. To this end we used the BIC to select the best model and to ensure that it was not possible to fit the data with a single distribution.

Most of our results depend on a linear model of the firing-rate. But does this relation have to be linear (and not, for example, logarithmic or sigmoidal)? Firstly we note that linear functions have been extensively used to model the relation between the firing-rates of neurons and certain task features (e.g. **??**Pardo-Vázquez et al. (2008)). Yet it is possible for the relation not to be linear. Indeed, we consider the linear function as a first probable approximation. Hence, we also analyzed firing-rates with the mutual information technique that also decodes non-linear relationships. The results, however, were not much different than the linear model (data not shown). We finally decided on a linear approximation as this provides for a more intuitive interpretation.

The three mechanisms underlying the confidence x-shaped pattern that we have suggested, raise the question of whether PMv neurons change their firing-rate gradually, or whether they jump from a low to a high

activity state. This question, that has often raised concerning the decision neurons of the lateral intraparietal sulcus (LIP), has been bothering the scientific community for some time now (Horwitz and Newsome, 2001; Roitman and Shadlen, 2002a; Churchland et al., 2011; Ponce-Alvarez et al., 2012). Recently, Bollimunta et al. (2012) have provided reliable evidence for the hypothesis that LIP neurons display a gradual ramp. Although our analysis was aimed at differentiating single trial mechanisms, we did not address this issue. We do note that all three proposed mechanisms are compatible with both a gradual and an abrupt transition of states.

In these last two chapter we provided a computational framework with the support of neurophysiological evidence for decision confidence processes involved in post-decision wagering tasks. As outlined in sec.1.2.3 another widely used experimental procedure to measure confidence is the "uncertain option" task. In the next chapter we provide a biologically plausible model that can account for confidence-related decisions in the contex of the uncertain option task.

3.4   Methods

3.4.1   The Discrimination Task

Experiments were made using two male monkeys (Macaca mulatta). Animals (BM5, 8 kg; and BM6, 6 kg) were handled according to the standards of the European Union (86/609/EU), Spain (RD 1201/2005), and the Society for Neuroscience Policies and Use of Animals and Humans in Neuroscience Research. The experimental procedures were approved by the Bioethics Commission of the University of Santiago de Compostela (Spain).

The monkeys' heads were immobilized during the task and looked binocularly at a monitor screen placed 114 cm away from their eyes (1 cm subtended 0.5 to the eye). The room was isolated and soundproofed. Two circles (1° in diameter) were horizontally displayed 6° at the right and 6° at the left of the fixation point (a vertical line; 0.5° length, 0.02° wide) displayed in the screen center. The monkeys used

right and left circles to signal with an eye movement the orientation of visual stimuli to the right and to the left, respectively. Orientation Discriminations Task: the monkeys were trained to discriminate up to their psychophysical thresholds in the visual discrimination task sketched in Figure 3.1A (training lasted for approx. 11 months). The stimuli were presented in the center of the monitor screen and eye movements larger than 2.5° aborted the task. The orientation discrimination task was a two-interval, two-alternative forced-choice task. A masking white noise signaled the beginning of the trial and then the fixation target (FT) appeared in the center of the screen (Fig. 3.1A). The monkey was required to fixate the FT. If fixation was maintained for 100 ms, the FT disappeared, and, after a variable pre-stimulus delay (100 300 ms), two stimuli (S1 and S2), each of 500 ms duration, were presented in sequence, with a fixed inter-stimulus interval (1 s). At the end of the second stimulus, the subject made a saccadic eye movement, in a 1200 ms time window, to one of the two circles, indicating whether the orientation of the second stimulus was clockwise or counterclockwise to the first. Trials lasted approx. 3.5 s separated by a variable intertrial interval (1.53 s). Fifty milliseconds after the correct response, a drop of liquid was delivered as a reward. A modulation of the masking noise signaled the errors; the modulation started 50 ms after the incorrect response and lasted for 75 ms.

Monkeys weights were measured daily to control hydration, and once a week the animals had access to water ad libitum. The level of training was assessed by the psychometric functions. Once trained, the monkeys performed around 1000 trials per day. The lines were stationary, subtending 8° length and 0.15° wide. Three different S1 orientations were used for each monkey during the recordings: 87°, 90°, and 93° (BM5) and 84°, 90° and 96° (BM6); all angles referred to the horizontal axis. Different S2, eight per S1, were presented, four clockwise and four counterclockwise to S1 in steps of 1° (BM5) and 2° (BM6). More details can be found in Pardo-Vázquez et al. (2008).

3.4.2   Recordings

Neuronal population: extracellular single-unit activity was recorded with tungsten micro-electrodes (epoxylite insulation, 1.5-3.5M, catalog # UEWMGCLMDNNF; FHC) in the posterior bank of the ventral arm of the sulcus arcuatus and adjacent surface in the ventral premotor cortex in the four hemispheres of the two monkeys (see Pardo-Vázquez et al. (2008), for a detailed description of the recording sites). In this work, we studied the responses of a subset (336) of the recorded neurons. This subset was selected with a ROC analysis of firing-rate respect to the choice (see Pardo-Vázquez et al. (2008) for details).

3.4.3   Data Analysis

All analyses were performed using custom-made programs in Matlab. Unless noted otherwise, all statistical analyses were applied to the firing-rates of single neurons during the 500 ms preceding the saccade. In fact, the second stimulus was presented during this period, and therefore the decision-making process was expected to take place during this time window. Our aim was to found any existing neurons whose activity relates to:

1. Difficulty of the task: Neurons whose mean firing-rate for the correct trials increases linearly with the difficulty. When plotted against the TRO the activity of these neurons shows the v-pattern.

2. Confidence: Confidence measures present a characteristic x-shaped pattern when plotted against the signed difficulty of the task (Vickers, 1979b; Kepecs and Mainen, 2012; Kepecs et al., 2008). Therefore, when a difficulty neuron had enough error trials to be analyzed and when its firing-rate in error trials showed a mirror modulation of the difficulty respect to correct trials we considered it a confidence neuron (i.e., neurons that are a subset of the difficulty neurons)

In order to accomplish this we used a linear regression analysis (LM)(Draper and Smith, 1966). Of course, linearity is only one of numerous possible

encoding mechanisms, even when we take only those concerning firing-rates into consideration. We decided on this for the sake of simplicity.

3.4.4 Error Trials

In order to identify confidence neurons we independently applied the LM to correct and error trials. Unfortunately, the number of error trials was not enough to analyze all the neurons. As experiments were done using animals that were awake it was very difficult to record single neurons over a long period. We recorded approximately 10 trials per monkey and stimulus conditions (i.e., orientation of S1, S2). Hence only few error trials were recorded under easier conditions ($|S1-S2| = 3, 4$). In the end only 124 neurons had at least one error for difficult categories ($|S1-S2| = 1, 2$). Therefore, we were only able to run the LM methods (described below) on this subset.

3.4.5 Linear Analysis

The firing-rate (FR) of the last 500 ms before the saccade was computed by averaging the spike count in a sliding window of 100 ms slided with a step of 20 ms. In this way we got for each trial and each neuron a time series $r(t)$ of the firing-rate, where t is time discretized in 25 time bins.

3.4.6 Difficulty Neurons

To individuate the difficulty neurons (v-shaped modulation) the following LM analysis was used, $r(t) = d_1 |S1 - S2| + d_2(t)$, where d_1, d_2 are the parameters to be fitted. Activity was considered linearly dependent on the difficulty, $|S1 - S2|$, if the coefficient d_1 was different from 0 ($p < 0.05$, t-statistics) and no sign switch occurred during a time interval (T) with a length of at least four consecutive bins (140 ms). Details as to how T was chosen are given below.

3.4.7 Confidence Neurons

To select the confidence neurons we applied the above LM analysis to both correct and error trials, but we kept the two analyses separate. Therefore, the pairs of parameters d_1, d_2 are different for error (d_1^e, d_2^e) and correct (d_1^c, d_2^c) trials. We looked for neurons whose firing-rate presented a mirror modulation of difficulty in error trials compared to correct trials, hence we considered a neuron as encoding confidence if the sign of d_1^c was the opposite of that of d_1^e. Due to the low number of error trials, it was only possible to analyze fifteen of the 49 difficulty neurons; only five satisfied all of the constraints described above. In order to produce fig.3.3 the firing-rate of each neuron was normalized to its maximum value and then the activity of all neurons was aaveraged together.

3.4.8 Minimal Time Window (T)

In order to find the minimum length of T we proceeded as follows. Given that linear regression has a p-value cutoff at 0.05 in each bin there is a probability of 0.05 to get a false positive. We wanted to know what the probability P_n is of getting n consecutive false positives. We then selected n such that $P_n < 0.05$.

In order to calculate P_n we proceeded as follows. A statistical test with $p < \alpha = 0.05$ applied to a time series produced a time series of significant and non-significant bins. The vector X representing the time series is generated by:

$$X = \{x_1, \ldots, x_{25} \,|\, x_i \in [0, 1]\,, P(x_i = 1) = \alpha, P(x_i = 0) = 1 - \alpha\}$$

where x_i takes value 1 when the ith bin is significant by chance and 0 otherwise. X has length 25 since our time series $(r(t))$ has 25 bins. We generated 10^6 vectors with this procedure and then evaluated the probability P_n of having n consecutive ones. P_n is thus the probability that an i exists such that $x_i + \ldots + x_{i+(n-1)} = n$. Since we ran the test on a large number of neurons we corrected P_n for the family wise

error-rate: $P_n^N = (1 - (1 - P_n)^N)$, where N is the number of neurons (336). Then we found n such that $P_n < 0.05$. We found that the minimum number of consecutive bins needed to get a significant result was $n = 4$. The applied method gave the same results of the more common Bonferroni correction.

3.4.9   Mechanism for the Difficulty Neurons

We individuated three possible neural mechanisms responsible for the above mentioned modulation of the difficulty neurons. A simplified representation of these mechanisms is presented in Fig. 3.4. In order to understand which difficulty neuron belongs to each of the three categories, we applied two methods:

1. In order to find neurons that switch states with a timing dependent on the difficulty (i.e., signaled the decision with a change in activity), we used the Hidden Markov Model (HMM) analysis (Rabiner, 1989). Indeed, the HMM was able to cluster the spiking activity of individual neurons into periods of 'stationary' activity (the states) within a single trial. Hence the switch time between states could be estimated.

2. In order to find neurons whose activity after the change encoded the difficulty we calculated the correlation between the mean activity and the difficulty of the task. The mean activity was calculated in the time window starting at the time bin where the 90% of the trials had passed from one state to the other and ending at the last significant time bin marked by the LM.

3. In order to find the neurons whose activity could be explained as a compound of high and low firing-rate states we fitted (χ^2 goodness-of-fit test) the firing-rate distribution to the average of two gamma distribution functions.

3.4.10 Hidden Markov Model

To analyze the single-trial activity of the recorded neurons we used the HMM that clusters the spiking activity of individual neurons into periods of stationary activity within a single trial. The HMM technique has been successfully applied to characterize the single-trial activity of cortical neuronal ensembles during movement with holding and preparation (Seidemann et al., 1996; Kemere et al., 2008), taste processing (Jones et al., 2007), and perceptual decision making (Ponce-Alvarez et al., 2012). Here, we briefly review some aspects of the HMM analysis; more details about the algorithms can be found in previous works (Seidemann et al., 1996; Jones et al., 2007; Ponce-Alvarez et al., 2012).

Within the HMM, the activity of a recorded neuron at time t is assumed to be in one of a (predetermined) number (Q) of hidden firing-rate states. In each state q, the discharge of a neuron is assumed to be a Poisson process of intensity λ_q, which defines the instantaneous firing probability E_q, i.e., the probability of firing a spike within one time bin, equal to 2ms throughout this study. States are said to be hidden because they are not directly measured; instead, we observe the stochastic realizations of the state-dependent Poisson process (observation sequences). The state variable changes from state i to state j with fixed probabilities that defined a transition matrix A, given by $A_{ij} = P(q_{t+1} = j | q_t = i)$, where q_t is the state at time t and $i, j \in \{1, \ldots, Q\}$. The entire process is a Markov chain: the transition probabilities A_{ij} are independent of time, i.e., they depend only on the identities of states i and j, which means that the state sequence at time t only depends on the state at time $t-1$. In summary, for a single neuron the HMM is fully characterized by the spike-emission probabilities (E) and the transition matrix (A). These model parameters are estimated from the data, using a likelihood expectation-maximization algorithm (Seidemann et al., 1996; Jones et al., 2007; Ponce-Alvarez et al., 2012).

Briefly explained, the procedure starts with random values for E and A and re-estimates the parameters to maximize the probability of observing the data given the model. After optimization of the model parameters, the Viterbi algorithm is used to find the most likely sequence of hidden states given, for each single trial, the model and the

observation sequence (Ponce-Alvarez et al., 2012). In the present study we used the HMM to detect the transitions between a state of low and a state of high activity. For this reason, the number of states was set to $Q = 2$. For each neuron, the data was divided into two subsets, composed of trials corresponding to each behavioral response (left or right). For each subset, a HMM was estimated using the activity of 80% of the trials (randomly selected) during the period within the last 500 ms before the saccade. After optimization the most likely state sequence was stored for all trials.

Unfortunately, a HMM analysis was not reliable for all the neurons. We only considered the HMM reliable if a) The duration of both states was at least 25 ms. (i.e., we do not take into account states with very brief duration) b) The number of state-switches per trial was three or less; or (i.e., we do not take into account bursting neurons) c) At least five of both the left and right oriented trials had a state-switch (i.e., we want neurons with 2 different states). We found 45 confidence neurons (out of 49) whose HMM was interpretable. For this subset, we wanted to distinguish between the three v-shaped mechanisms, to do so we analyzed the state-switch time. For each trial the HMM gave the time in which it changed from a low to a high state (or vice versa).

3.4.11 Bimodality vs Unimodality

Our aim was to investigate whether the firing-rate distribution of neurons during correct trials was better described using a bimodal than a unimodal function distribution. The procedure we applied was the following:

1. Trials were divided into two sets, depending on their behavioral responses (left or right). We calculated the average firing-rate for each trial in a 200 ms time window that ended at the time of subject's behavioral response. We called the empirical distribution functions of the average firing-rate $F(\nu_L)$ and $F(\nu_R)$ respectively. The distributions were fitted with a function B, which was the average of two gamma distributions: $B =$

$(F_\Gamma(x; \kappa_1, \mu_1) + F_\Gamma(x; \kappa_2, \mu_2))/2$. Gamma distribution was chosen because is one of the most general function distributions with positive support. The gamma distribution is given by:

$$F_\Gamma(x; \kappa, \mu) = \frac{1}{(\mu/\kappa)^\kappa} \frac{1}{\Gamma(\kappa)} x^{\kappa-1} e^{(-\frac{x\kappa}{\mu})}$$

for $x, \kappa, \mu > 0$; Γ is the gamma function; κ is the shape parameter and μ is the mean.

2. We looked for best fit (in terms of χ^2 goodness-of-fit, $p < 0.05$) using the following parameter space: $\mu_d \in [min(F(\nu_d)), max(F(\nu_d))], \kappa_d \in [0.1, 10]$, for $d \in (1, 2)$.

   In order to apply the goodness-of-fit test, we discretized the firing-rate distribution in bins of 5 Hz. The χ^2 goodness-of-fit probability test is valid under the assumption that the number of events in each bin is greater than five. Whenever this condition was not satisfied we enlarged the bin on the right until the event count was at least five.

3. Finally, in order to verify whether a bimodal model explains the data better than a model with just one mean, we fitted the data to a single Gamma distribution and then we compared the two models using the Bayesian Information Criterion (BIC) (Schwarz, 1978; Akaike, 1974), which is given by:

$$BIC = -2 \cdot \ln L + p \ln(T)$$

where L is the maximized value of the likelihood function for the estimated model; p the number of free parameters of the model (2 or 4); and T the length of the observation data (the number of bins). This means that the BIC method penalized the model likelihood by a measure of its complexity (i.e., the number of free parameters). The single mode model has two free parameters, while the bimodal model has four free parameters. Therefore, in order to have a better score in the BIC, the higher complexity due to the second mode should really be well balanced by a better ability to explain the data. Hence we

considered that a neuron has a binary response if the BIC favored the
model with two modes.

The Uncertain Model

The results of this chapter are included in Insabato, A. et al., Flexibility in decision making: neural mechanism for confidence-driven multiple choices. Submitted.

4.1 Introduction

In the preavious two chapters we prepared the building blocks for developing a computational framework to understand confidence processes in post-decision wagering experiments and provided experimental evidence for the model proposed in chap.2. In this chapter we are going to propose a new model to account for behavioral and neurophysiological data in an uncertain-option task. We will then discuss the differences between the two models and the implication of this dichotomy for our general understanding of the phenomenon of decision confidence.

We recall here very briefly the two experimental setups used to record confidence with non human subjects: post-decision wagering and uncertain-option. In the first type of experiment, after decision, subjects bet on the outcome of their choice according to their confidence. In the latter subjects are given the choice in each trial between performing a decision-making task, which could lead to a reward if the answer is

correct, and selecting a sure but less valuable reward (for an exended discussion see sec.1.2.3 and Kepecs and Mainen (2012)).

Recently a notable evidence has been made available by the study of Kiani and Shadlen (2009). They recorded neural activity in LIP using a combination of the random dot motion (RDM) task with the uncertain-option task, allowing the subject to opt for a small but sure reward instead of deciding about the direction of motion. LIP neurons are known to receive input from MT neurons, that encode the motion energy (Britten et al., 1996b; Snowden et al., 1991a; Simoncelli and Heeger, 1998), and to fire according to subject decisions in motion dicrimination tasks both in binary (Roitman and Shadlen, 2002a) and multiple choices task (Churchland et al., 2008a). LIP neural activty underlying the decision process has been successfully accounted for by attractor neural networks (ANN) both for the binary (Wang, 2002b) and the multiple choices task (Albantakis and Deco, 2009b).

The task introduced by Kiani and Shadlen (2009) poses new theoretical questions since it requires, beside the usual decision, a dynamic online enconding of confidence, very rapid evaluation of the uncertainty in the decision and the late integration of a third alternative into the decision process.

In this chapter we want to produce a very simple model capable of solving the uncertain-option task by dynamically encoding decision confidence. In contrast to chap.2 we propose that in this case there is no need for a separated monitoring network and that confidence computations can occure solely into the decision making circuit. This provides a more minimalistic account of some confidence processes and enrich the actual discussion about confidence by separating confidence and monitoring. Moreover in this task the final decision depends on inputs arriving at different timing therefore our work sheds also some light on the dynamics of decision-making attractor networks in a scenario with asyncronous inputs.

We propose that confidence, measured in an uncertain-option task, is encoded in the state of the decision variables of a multiple choice decision making neural network model working in the multistable region of the attractors landscape. Our model is indeed able to reproduce both psychophysical and neurophysiological data of Kiani and

Shadlen (2009). This means that confidence can be an emerging phenomenon of decision-making networks and that a monitoring network is not strictly necessary. We describe an "undecided" region, in the state space, related to lack of confidence. The more time the system wonders in this region the higher the probability that the decision will be aborted. Furthermore we make several predictions analyzing decision times and the decision state before the overt commitment at the end of the task. We predict that a reaction time version of the task would report a bimodal distribution of reaction times (RT). Moreover we refute the classical criticisms to uncertain-option task (Shields et al., 1997; Smith et al., 2006) by predicting that the probability of choosing a safer option changes dramatically in early correct vs error trials, showing the signature of a confidence measurement. Last but not least our model shows that the subjects adopt a flexibile strategy and prefer to abort the decision in favor of a safer option at the expenses of performance. Our results show that this aversity to risk is maximal in the multistable regime of the model. This in turn predicts different results when a risk-seeking behavior is promoted.

4.2 Results

4.2.1 Confidence through Multiple Choice Mechanism

The main idea of our study is that simple confidence mechanisms (e.g., the one involved in uncertain-option tasks) can be accounted for in a multiple-choice decision-making network by encoding the confidence in the state of the decision variables. As evidence for such confidence processes, we referred to the findings of Kiani and Shadlen[2009], who, in seeking to distinguish high versus low confidence trials, modified the classical RDM task by introducing a sure option.

The study of Kiani and Shadlen (2009) was already described in chap.1 but we summarize it here fot the comfort of the reader. The typical sequence of events in a trial is represented in fig.1a.
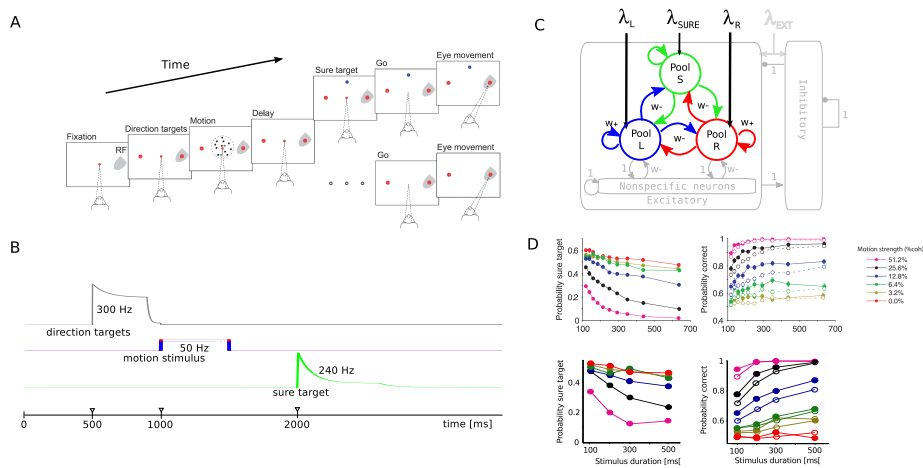
Figure 4.1: Task description, network structure, stimulation protocol and psychophysics. (A) Experimental task procedure of Kiani and Shadlen (2009). (B) Stimulation protocol for populations due to direction targets presentation, motion stimulus and "sure" target. Both direction selective pools receive the same targets input while the motion input can be different across pools in order to represent different values of motion coherence. Pool S only receives input due to the presentation of "sure" target. (C) Network scheme representing the connection structure between pools and inputs. (D) Psychohysics results of the model (bottom row) are compared to experimental data (top row). The model reproduces qualitatively the effect of both stimulus duration and coherence on both probability of correct responces and probability of choosing the "sure" option. Moreover the model reproduces the increased $P(correct)$ when the "sure" option was waived (filled circles).

The subject is required to maintain fixation in the center of the screen in order to start a trial, then the targets appear on two opposite positions corresponding to the two possible decisions about the direction of motion. The motion stimulus is then showed for a variable duration followed by a variable delay period during which the monkey must hold the decision in memory. The end of the delay period is marked by a go signal that triggers the saccade of the subject toward the decided motion direction. In a random half of the trials, during the delay period, the monkey was presented a third "sure" target. The

election of this target corresponds to waiving the decision but was always rewarded with a smaller amount of liquid than the direction targets. The proportion of trials on which the decision was waived can be considered a measure of average decision confidence in a block of trials. Indeed Kiani and Shadlen (2009) found that the probability of choosing the "sure" target ($P(S)$) decreases both for longer stimulus durations (which should correspond to more certain conditions) and for higher motion coherences (i.e., easier stimuli) (Fig. 1D). They also found that the probability of a correct choice, when the "sure" target was shown but waived was higher than in trials without "sure" target presentation, which is an important finding to argument that the subjects are actually monitoring their cognitive processes and not only learning an appropriate stimulus response association. In the context of this task the probability of a high confidence report and the probability of choosing either "left" or "right", which is equal to $1P(S)$, are the same. The probability of a high confidence report can in turn be considered as reflecting a continuous value of confidence. How the P(S) can reflect the confidence, or the lack of it, in a decision is an undoubtedly interesting (and complicated) question but we aren't going to address it here and will therefore use the terms confidence and probability of "sure" response without any distinction.

The neural recordings of single neurons in LIP showed, beside the well known accumulation of evidence and the differential activity due to decision [Roitman and Shadlen,2002;Gold and Shadlen,2007;Bollimunta and Ditterich,2011], an intermediate level of activity associated with choice of the "sure" target (Fig. 2). One of the aims of this article is to demonstrate that a network designed for multiple-choices perceptual decision-making is able to show a behavior attributable to confidence processing. Indeed, when fed with an input that encodes the value of the "sure option" this network can take charge of the evaluation and use of the confidence information about the perceptual decision. To this aim and in order to account for data of Kiani and Shadlen[2009], we build a network with integrate-and-fire neurons using different synaptic dynamics for AMPA, NMDA and GABA receptors. Figure 1 graphically summarizes network details. The network has one population of neurons for each decision to be taken. Population R and L, for right and left targets, and population S for

the "sure" target. The response of MT neurons is almost linearly related to the coherence level of random dots motion (Wang, 2002b). Consequently we increased the difference between inputs to population R and population L, according to the prevalent direction of dots motion. Each pool receives also an input due to targets. Recently Rorie etal.[2010] found that the firing rate of LIP neurons during target presentation is proportional to the reward associated with the target, therefore we set the target related input to a higher value for high stakes pools (R,L) respect to the "sure" pool (S). Connections strength was designed in order to implement a competition mechanism between pools [Wang,2002;Albantakis and Deco,2009]. Differential input ($\Delta\lambda = \lambda_L\lambda_R > 0$) to the pools can bias the decision-making, but given the intrinsic stochasticity of the dynamics, an incorrect decision may be taken even when input currents are very different.

Our model implements a simple multiple-choices decision-making mechanism [Albantakis and Deco,2009]: if three synchronous inputs (reflecting three possible directions of dots motion) are given to the network it would just take a decision about which of the three inputs is the largest. However we used an asynchronous stimulation protocol like the one used by Kiani and Shadlen[2009], which has not yet been studied with attractors networks, that segregates the decision process in two stages: a first decision between two alternatives and a second decision between the chosen one and a third alternative. In this context the decisions don't produce a behavioral response but are stored in memory for a later use. As shown in Fig. 2, the model receives input from the RDM stimulus once the targets are presented. This input triggers a competition between the two pools, L and R, soon one of them prevails, increasing the firing rate. During the delay period its firing rate remains high, keeping trace of the decision [Brunel and Wang,2001]. When the input to pool S is turned on, the high activity of the winning pool is inhibiting, through inhibitory neurons (see fig.1), the other pools preventing population S neurons from increasing their firing rates. On some trials the integration of evidence takes longer and eventually the stimulus ends while R and L have still similar activity and no decision has been taken. On these trials typically pool S receives less inhibition and therefore can increase the firing rate and win the competition against R and L.

We highlight the fact the model was not designed for decision confidence representation. The underlying mechanism is a simple decision-making process and the input was constructed to mimic MT response to visual stimuli. The key ingredient in order to convert the network to a confidence evaluation process is the input. Indeed the input to pool S does not contain information about the correct response but rather information about the value of the "sure" alternative (as experimentally shown by [Rorie etal.,2010]). In this scenario the confidence-related behavior arises naturally given the interaction between network intrinsic properties and input structure.

4.2.2   Psychophysics and Neurophysiology of Decision Confidence

Our model can explain both the psychophysical and the neurophysiological data (fig. 4.1D and 4.2) as activity of the multiple choices decision-making network decribed in the preavious paragraph. In addition to the details explain above, a key ingredient of the model for reproducing the data of Kiani and Shadlen (2009) is an attractors landscape where three attractors are stable at same time (also known as multistable regime). In general decision-making ANNs can have several attractors and the parameters (e.g. connectivity, inputs, etc.) can manipulate the attractors landscape. Usually decision-making tasks only require one attractor for each alternative (Wang, 2002b; Albantakis and Deco, 2009b; Marti et al., 2006). Conversely we hypothesized that a third non-decision attractor, where both decision pools have comparable firing rate, could be a good representation for a low confidence state.

We show in fig. 4.3E the bifurcation diagram of the system (a mean-field equivalent of the spiking network). This diagram shows the firing rate of the decision pools when the system converged to an attractor for each attractor and as a function of the parameter λ, the common input to the decision pools (see sec.4.4).

As can be observed the variation of λ induces different dynamical regimes of the model. At $\lambda = 0$ Hz there is a "spontaneous attractor", where both pools are in a resting state around 2 Hz, and two
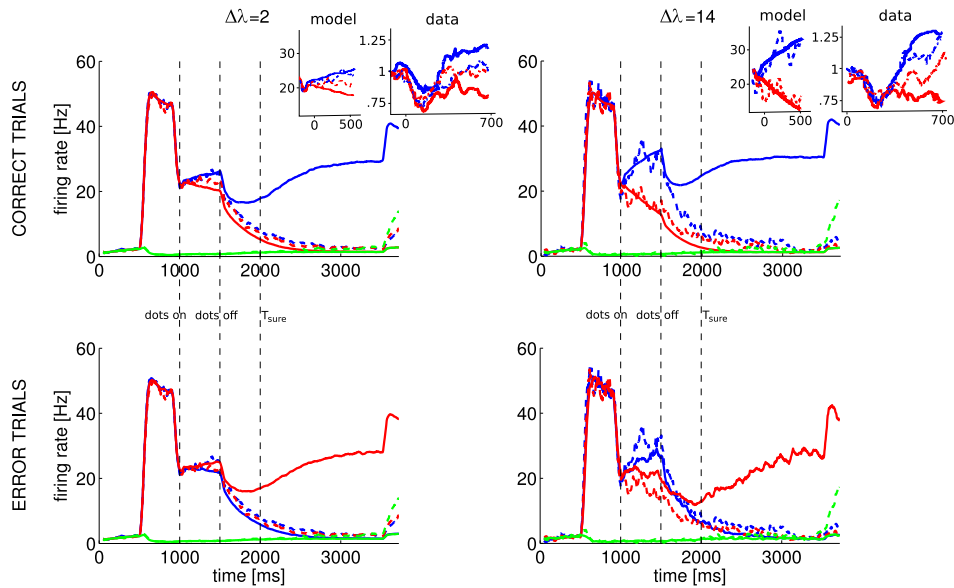
Figure 4.2: Model dynamics. Red/blue/green lines represent R/L/S pool firing rates. Continuous lines are average firing rates for trial when one of the two direction option was chosen. Dashed lines are average rates of "unsure" trials (i.e. pool S won). The network shows different dynamics depending on the winning pool. When the final decision was to opt out for the sure bet target the firing rates of the decision pools remain similar during input presentation, meaning that a decision has not been reached. This behavior closely resambles the one observed in LIP by **?** as can be observed in the insets of the top row showing a magnification of the stimulus period and a comparison with experimental data (showing normalized firing rate). Kiani and Shadlen (2009) don't show error trials firing rates, so we are not able to compare them. The x axis in the insets represents time from stimulation.

decision attractors. This decision attractors in absence of input are key for keeping in memory the decision during the delay period. For values grater than 0 the spontaneous attractor destabilizes and only the decision attractors remain stable. At $\lambda = 20$ Hz another saddle-node bifurcation occurs and a stable "mixed" state appears, where both pools fire at high rate. This multistable region extends until $\lambda = 55$ Hz, where due to a subcritical pitchfork bifurcation the decision attractors disappear and only the "mixed" state remains stable.
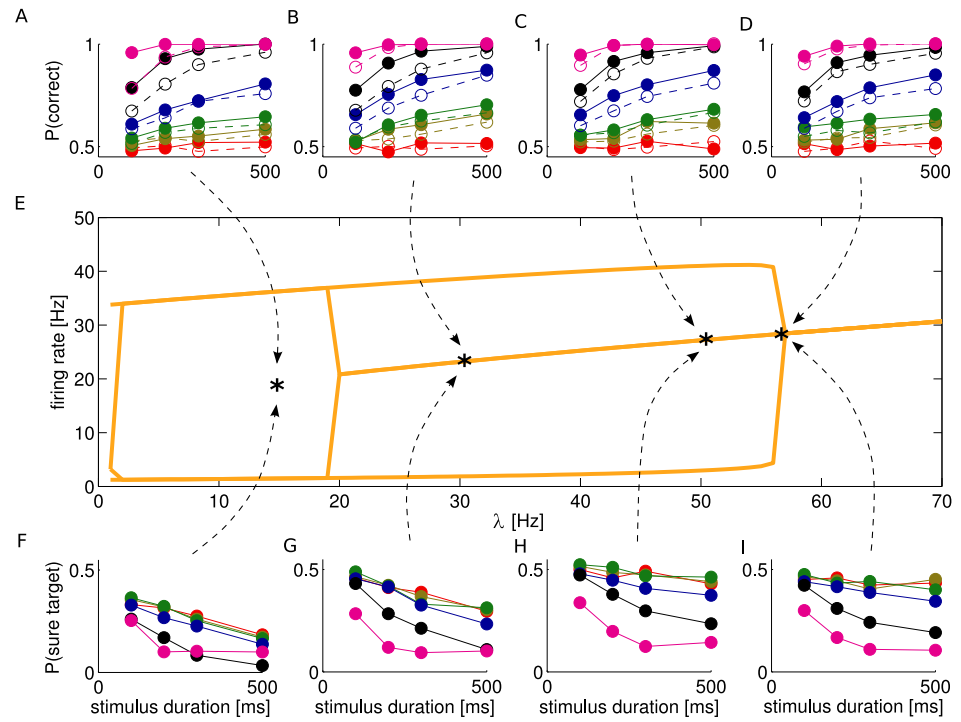
Figure 4.3: Bifurcation diagram and psychophysics measures in different regions. Proximity with second bifurcation increases flexibility and optimizes reward amount.

Panels A-D show the $P(correct)$ for 4 sample points in different regions of the bifurcation diagram (marked by stars). For each point we adjusted the parameter $\Delta\lambda$ in order to get a good match with experimentally measured $P(correct)$. Panels F-I show the $P(S)$ for the same points. As can be seen the model is able to reproduce qualitatively the performance of monkeys in the entire region considered (we didn't take into account the region to the right of the pitchfork since it is clear that decision-making is impossible with just one attractor). Interestingly the $P(S)$ changes qualitatively when moving through the different regimes of the network. In the bistable region (light blue, panel F) the $P(S)$ is less sensitive to changes in stimulus coherence (i.e. $\Delta\lambda$) and rapidly decreases as a function of stimulus duration for low evidence trials (red curve), while experimental data show a very

flat curve. When moving towards the third bifurcation this effect is gradually reduced and a good qualitative fit of experimental data can be observed in the light red region, just before the third bifurcation. Changes in the input to the S pool have the main effect of shifting the curves up or down but don't imply a qualitative change (data not shown).

The model also reproduces the increased performance in trials when the "sure" target was shown but not chosen (filled circles in fig. 4.1D and 4.3A-D). While the $P(correct)$ and $P(S)$ depend on network parameters and required some tuning in order to be adjusted to experimental data (see text below and fig. 4.3), this last feature was not obtained by tuning the parameters and is therefore a genuine feature of the model that is confirmed by the experimental evidence.

Simulation results also reproduce main neurophysiological effects observed in LIP. In fig. 4.2 we show the average population firing rates for the decision pools (L,R,and S) for easy and difficult trials and for correct and error responses. Firing rates show the typical abrupt response to target presentation and the rapid decay just before stimulus presentation; then, in trials when the sure target is not chosen (solid lines), after the stimulus is presented, neurons show the ramping activity commonly described in LIP (Mazurek et al., 2003). During the delay period the firing rate drops down but mantain a sustained activity for holding in memory the decision. Finally the saccadic eye movemen is signaled by the increase in firing rate at the end of the trial. In unsure trials, when S pool wins (dashed lines), due to presence of the stable "mixed" attractor, pools R and L show an intermediate level of activity during stimulus presentation. The level of activity associated with unsure trials stays right in the middle of solid traces for difficult trials while in easy trials the firing rates of R and L are more separated, albeit remaining less separated than solid traces. This same feature appears in the data recorded by Kiani and Shadlen (2009) (shown in the inset for comparison).

4.2.3 Confidence is Related to the State of Decision Neurons

Detailed models are useful since they make possible to explain cognitive high level phenomena in terms of the underlying neural dynamics. We want to exploit this possibility and link the confidence level in a decision with the state of the neural system implementing the decision mechanism. Therefore here we want to track the confidence during the triansient period before reaching a decided state. To this end we can observe the decision process in the phase plane of firing rates ν_R and ν_L, as depicted in fig. 4.4A.

In contrast to diffution-like models (Kiani and Shadlen, 2009; Ratcliff and Rouder, 1998), that use an accumulation to threshold mechanism to determine the decision, in ANN models whenever the system enters into the basin of attraction of one of the decision attractors it will remain there and hence that decision will be taken[1] (Wang, 2012; Lo and Wang, 2006). In fig. 4.4 we plot the attractors and their basins in the phase plane of the system (the mean-field reduction of the spiking network) in the best fitting region of the bifurcation diagram ($\lambda = 50$ Hz). Dashed curves mark the basins of decision attractors when the stimulus is on. The region of the plane laying between these two zones is the basin of the "mixed" attractor (black dot on the diagonal) and represents an "undecided" state. During the delay period the stimulus is off and the attractors landscape is modified as shown by gray dots (dotted lines mark the basins under this condition). The trace in the plane, shows an example trial of the spiking network: when the motion

---

[1]We want to make clear the distinction between the decision mechanism intrinsic to the model and the mechanism used to read out that decision. If we imagine the condition of a decision when the stimulus doesn't provide any evidence it is particularly clear that a diffusion process won't reach a stable state and will change drift continuously upwords or downwards. On the contrary an ANN will reach a decision attractor and stay there (although change of mind are possible as shown by Albantakis and Deco (2012)). Therefore we could say that the ANN reaches a decision state, whenever it enters into the attractor. This state can then be read out by checking whether the firing rate of a population exceeed a given threshold (a commonly used method also with diffusion models) or by other mechanisms, as mentioned in sec.1.2.2. Diffusion models instead don't have a decision state, therefore their decision can only be determined when reading out their activity.
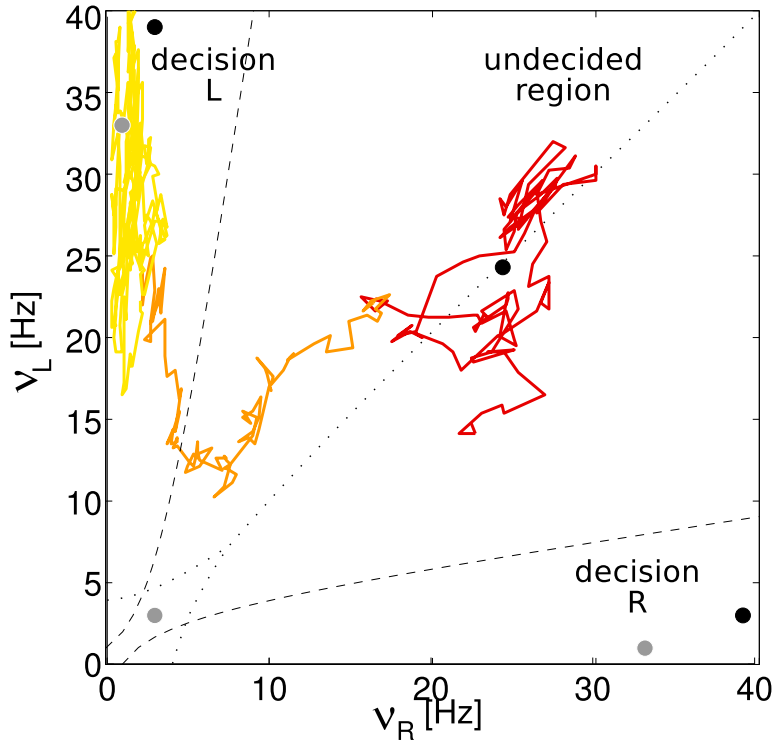
Figure 4.4: Basins of attraction and example dynamics. The figure represents an example trial (color represent three different moments of the trial: red, stimulus; orange, delay period until "sure" target onset; yellow, delay period until saccade) in the state space of firing rates of the two decision populations ν_R and ν_L. Dots represent the attractors position during motion stimulus (black) and in the delay period (gray). Dashed/dotted lines mark the boundaries of basins of attraction during the stimulus/delay period.

stimulus is turned on both populations present relatively high and comparable firing rate due to the preavious response to the direction targets. Then the system starts to move around the diagonal (red part of the trace), i.e. firing rates remain similar, due to the presence of the undecided attractor. Once the stimulus is over (orange part of the trace) the attractors landscape changes and therefore the trace moves towards one of the two sides getting into a decision region, i.e. one of the two pools increases the firing rate while the other decreases it. Our hypothesis is that the distance of the system from the decision region

determines the confidence and in turn the "sure" choices. In order to represent the overall state of the system under certain conditions we show the distribution of states of the full spiking network at the time of "sure" target onset, a snapshot of the decision process (fig. 4.5).
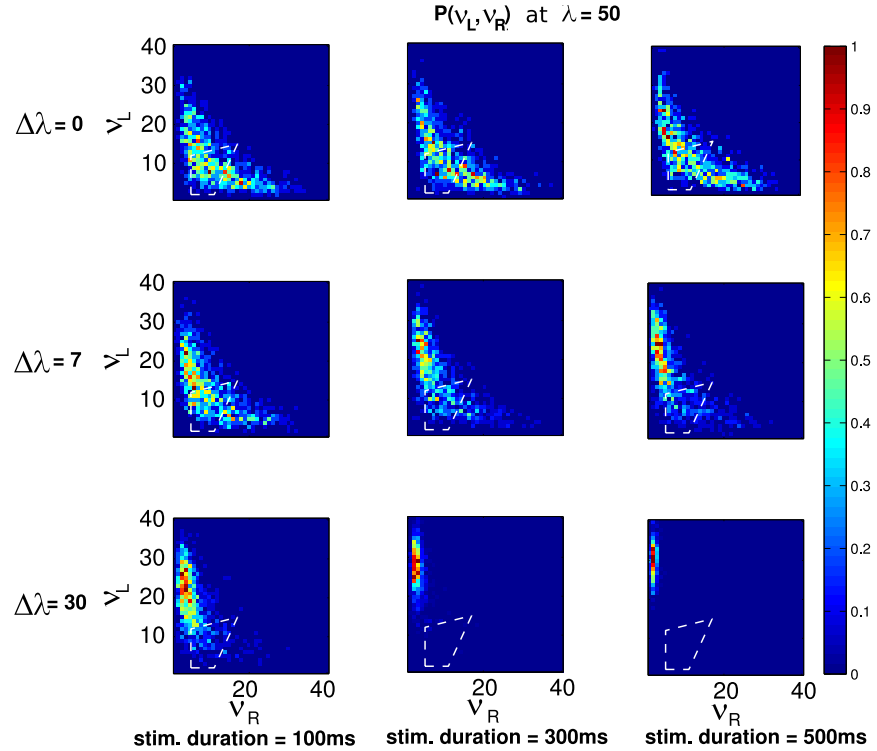
Figure 4.5: Distributions of firing rates. Each panels shows the distribution of firing rates at "sure" target onset for a given combination of stimulus evidence ($\Delta\lambda$) and stimulus duration. The dashed polygon is the convex hull of points $P(S|\nu_L, \nu_R) > 0.6$ as explained next and in Methods section.

The panels of fig. 4.5 represent different conditions: rows from top to bottom are distributions for increasing values of stimulus coherence, while columns from left to rigth represent increasing durations of the stimulus. When there is very low evidence (or no evidence at all, like in the first row) the distribution of ν_A, ν_B is quite symmetric about the diagonal and the duration of the stimulus has very low or no effect. In the second row we can observe that a moderate level of coherence

already skews the distribution towards one decision area and that the increasing duration of the stimulus makes the skewness even stronger. On the bottom row (very high stimulus coherence), for a short duration, the distribution appears markedly skewed towards the decision area and we can observe that for longer stimulus durations the distribution of system states moves even further into the decision region concentrating towards the decision attractor.

It is plausible to hypothesize that a prolonged stay in an undecided region is associated with lack of confidence. Indeed we found that the probability of choosing the sure option is very high when the firing rate of "left" and "right" neurons are similar and that it decreases very rapidly when those firing rates start to diverge, entering in a decision zone (see fig. 4.6).

Fig.4.6 shows the conditional probability of the "sure" response respect to the value of the firing rate of decision pools (ν_L/ν_R). It's evident that moving away from the diagonal the confidence increases (i.e. $P(S)$ decreases). In addition, the lower the activity of the two pools the lower the confidence. As a consequence, if we define the undecided region to be that portion of the plane where the $P(S|\nu_L, \nu_R) \geq 0.6$, then the time that the system spend in the undecided region depends on the duration of the stimulus and on the coherence, the two variables that manipulate the $P(S)$ (fig. 4.7). When the evidence for the decision is low (e.g. stimulus duration is 100 ms and $\Delta\lambda = 0$ the system spends more time in the undecided region and the decides more frequently to opt for the "sure" option compared to easier trials.

It is interesting that $P(S|\nu_l, \nu_R)$ does not change when the stimulus parameters (average intensity λ, duration and coherence or $\Delta\lambda$) are changed (see sec.4.A) What is changed manipulating the average intensity of the input (λ) or the amount of evidence ($\Delta\lambda$ and stimulus duration) is the distribution of firing rates in the ν_A, ν_B plane (as summarized in figs. 4.4 and 4.8) but not the $P(S|\nu_L, \nu_R)$.

Indeed, observing figures 4.4 and 4.8 we can see stimulus features dramatically change the probability of the system of being in a given position in the plane. For reference we plotted a dashed polygon in fig. 4.4 and 4.8 showing the convex hull of points with $P(S|\nu_L, \nu_R) = 0.6$.

Following this reasoning we can also explain why the network present
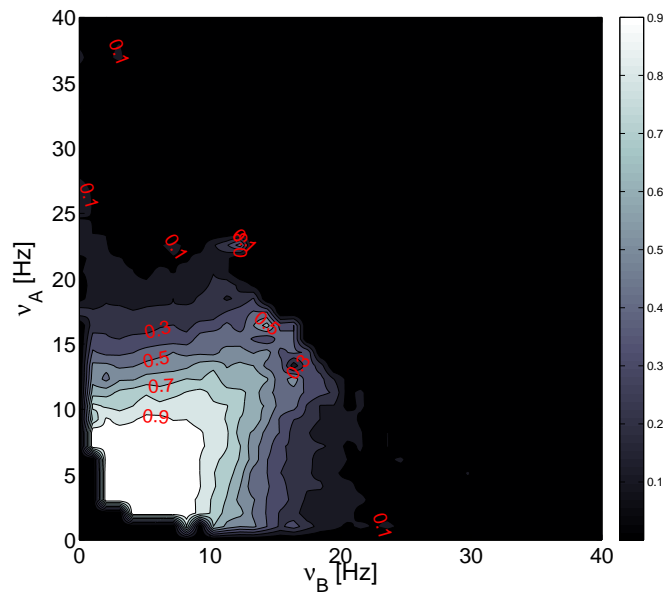
Figure 4.6: Conditional probability of choosing the "sure" option given the state of decision neurons. Each point of the plane represents a possible state of the system as characterized by the firing rate of decision pools (ν_L/ν_R). The color represents the probability that "sure" target is chosen, when the system in a given point of the space at the moment when the "sure" target is presented (see sec.4.4 for details). Contour lines pass through points with the same probability. It's evident that the $P(S|\nu_l, \nu_R)$ decreases rapidly when the system approaches the decision attractors.

a different behavior in the bistable regime compared to the multistable regime (see fig. 4.3F). If we look at the distribution of firing rates when λ is set to 15 Hz, i.e. in the bistable regime (fig.4.8), the most evident difference respect to network's behavior in the multistable regime is that the $P(S)$ decreases very rapidly when the duration of the stimulus increases. In particular for long durations the state of the network is concentrated in the decision regions where the $P(S)$ is very low (see fig. 4.6).
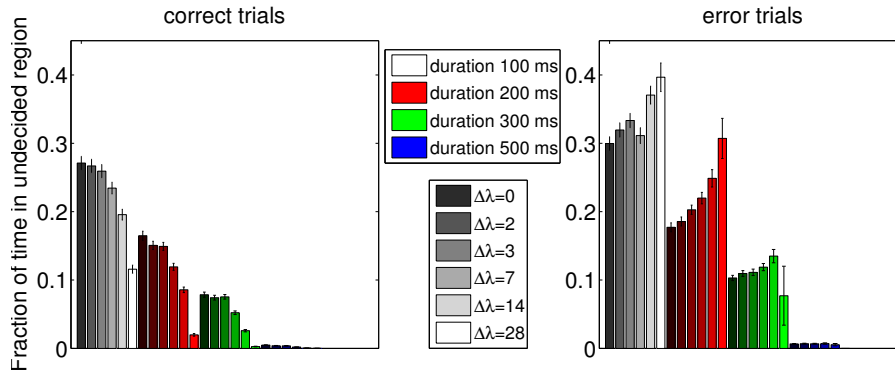
Figure 4.7: Time spent in the undecided region between stimulus onset and "sure" target onset for correct and error responces.

This explains why the probability of choosing the "sure" target decreases so rapidly as a function of stimulus duration in the bistable regime (see fig.4.3F).

Regarding the question of decision confidence encoding the model presented in this chapter makes a different proposal respect to other available models (see sec. 1.2.3 and chap.2). Indeed we can observe that the contour lines in fig. 4.6 form approximately square surfaces that strech slightly near the diagonal. This means that in general the network only takes into account the firing-rate of the winning pool for the confidence judgment. However when the firing-rates of the two decision pools are similar the network takes into account both pools. The confidence judgment seems therefore to be controlled by a distance function between the firing-rates of decision pools of the form $D_n = \sqrt[n]{|\nu_L^n - \nu_R^n|}$, for a large value of n. For $n = \infty$ the distance D_n becomes $d_n = \max\{\nu_l, \nu_r\}$ and the contour lines would be perfect squares. The classical "balance of evidence" suggested by Vickers (1979b) as the confidence code of race models is actually the simple distance between the two accumulators. Here we propose that the internal processing of the network apply a transformation to the simple difference between the two decision variables and produces a confidence, or actually a distrust encoding based on a non linear distance between the activity of decision pools.
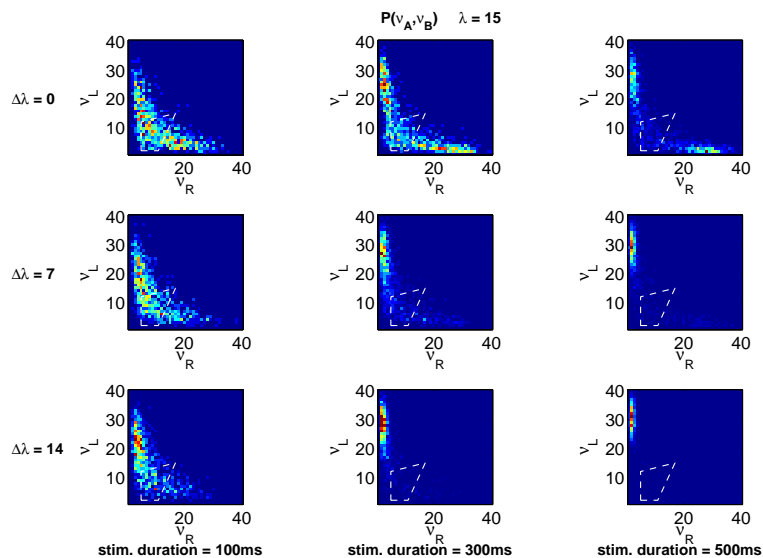
Figure 4.8: Distributions of firing rates in all trials for $\lambda = 15$ Hz. Panels A-I show the distribution of firing rates at T_{sure} onset for some combinations of stimulus evidence ($\Delta\lambda$) and stimulus duration. It can be observed that the high density in the zone with high $P(S)$ at very short durations rapidly vanishes for longer durations indicating that the uncertainty or lack of confidence in this system is mainly due to the dynamic effect of stimulus duration (not according to experimental results of Kiani and Shadlen (2009)).

4.2.4 Error Trials

Our trials last long (more than 2 s from stimulus on) due to the delay period but the decision process is typically quite rapid (about 1 s in our model for difficult conditions) and the decision is then stored in memory. If we could have access to the decision taken in the early stage of the trial we would expect a higher frequency of "sure" choices in error compared to correct trials, according to existing literature (see sec.1.2.3 and reference cited there). Unfortunately in the experiment this information in unavailable since on each trial only one responce is

given ("left", "right" or "sure"). Conversely in the model we are able to access at any time during a trial the state of the decision process (in the discussions we provide a possible paradigm to test our predictions). We call here the hypothetical outcome of the decison taken in the early stage of the trial "early correct" or "early error". The distribution of firing rates presents interesting differences when analysed separately for early correct (fig. 4.9) and early error (fig. 4.10) responces.

Figure 4.9: Distributions of firing rates in early correct trials. Panels A-I show the distribution of firing rates at T_{sure} onset for some combinations of stimulus evidence ($\Delta\lambda$) and stimulus duration.

While correct responces are distributed mainly in the decision region, error responces are more spread and occupy a larger area of the undecided region. At the same time on error trials the dynamics is slower and the system need more time to leave the undecided region, as shown by fig. 4.7. Indeed if we plot the proportion of "sure" choices separately for early correct and error trials a very interesing prediction arises. In fig. 4.11 we can see that while in correct trials the proportion of "sure" choices decreases when the task gets easier, in error trials it increases.
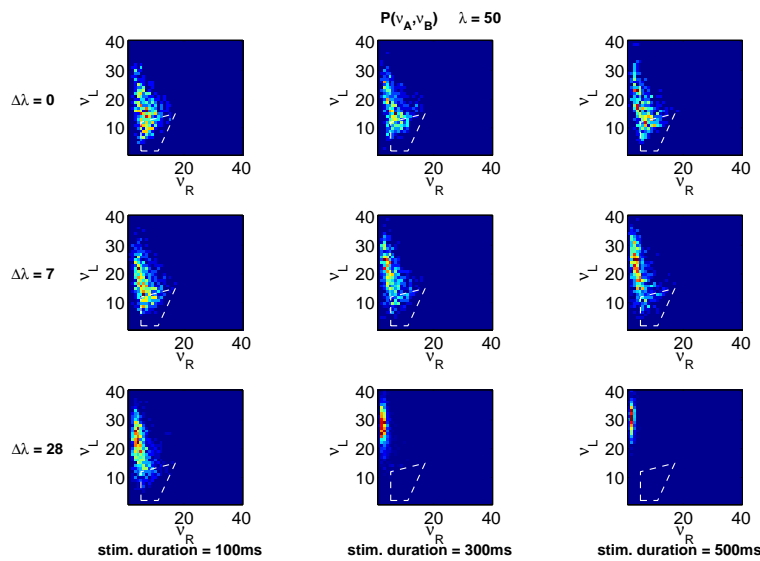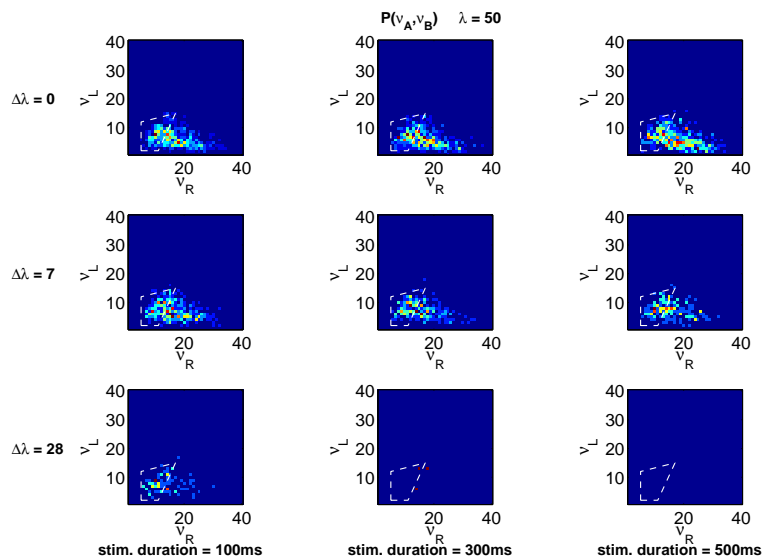
Figure 4.10: Ristributions of firing rates in early error trials. Panels A-I show the distribution of firing rates at T_{sure} onset for some combinations of stimulus evidence ($\Delta\lambda$) and stimulus duration.

This is exactly the pattern that one would expect from a confidence measure (Vickers, 1979b; Kepecs et al., 2008; Kepecs and Mainen, 2012). Moreover the increased $P(S)$ in early error trials explains the higher $P(correct)$ in trials when the sure option is shown but waived as a genuine confidence phenomenon invalidating the criticism against the unsure option task (Shields et al., 1997; Smith et al., 2006).

4.2.5 Confidence and Its Relationship to RTs

Exploiting the possibility to retrive information about the reaction time in the model we describe next the reaction times distributions and their relationship to confidence. What we refer to as reaction time is the time (from stimulus onset) that the model takes to integrate the stimulus signal and commit to a decision plus a non-decision time.
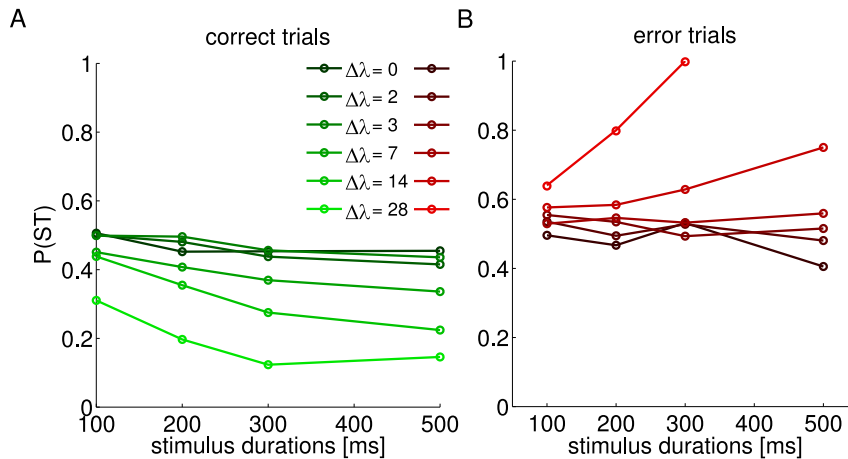
Figure 4.11: Probability of choosing the "sure target" as a function of stimulus duration shown separately for early correct and error trials. Different lines correspond to different values of $\Delta\lambda$ (i.e. stimulus coherence). This preditcion of the model shows that the probability of waiving the perceptual choice in this task really behaves like a lack of confidence measure, being higher in error trials and showing opposed modulation as a function of difficulty in correct versus error trials.

This is a standard way of accounting for reaction times in models of behavior since the non-decision time represent the sensory processing time and the time needed for motor exectution. Since our results show a qualitative effect of confidence on RTs and do not pretend to fit any behavioral data, the parameter of non-decision time in our analysis has been set to the arbitrary value of 80 ms but the results don't depend on this value. We consider here that a decision has been taken when the firing rate of one decision selective pool crosses a threshold fixed at 30 Hz (see sec.1.2.2 for a discussion about the decision criterion). Anyway we note that similar results are obtained in our model with a difference based mechanism. Indipendently on the decision rule another problem remains still open when facing with delayed response tasks. Indeed if after accomplishing the decision rule but before the "go" signal the accumulation process keeps on, a change of mind (Resulaj et al., 2009) is possible. Therefore it is not clear whether the decision is taken based on the winning accumulator when the decision criterion is met

(involving a separate working memory process) or based on the winning accumulator when the "go" signal is presented. We have considered that what matters is the winner when the threshold is crossed. In this scenario we report that in a $10 \pm 2\%$ of trials a change of mind, as described in Albantakis and Deco (2012), could occurs. This amount of changes does not affect our results.

We observe that in trials when "sure" target is not presented the reaction times have a bimodal distribution when the evidence for the decision is not very high. Fig.4.12 shows the histograms for a short and a long duration of the stimulus and for three increasing values of the coherence (i.e. $\Delta\lambda$). The distributions can be easily separated in a fast and a slow part, except for the bottom right panel. When the stimulus evidence increases (i.e. the coherence or the duration increases) the fast part of the distribution gets bigger while the slow part gets smaller. At some point, when evidence for the decision is very high, the two parts collapse into a unimodal left skewed distribution, as can be seen in the bottom right panel.

We note that the two parts of the distributions correspond to two different types of trials. On some trials the network reaches the decision criterion during the stimulus period and these trials result in fast responses. On the contrary when the a decision is not met during the stimulus period, once the stimulus is estiguished, the decision process can only be based on the memory trace of the stimulus (i.e. the self-sustained activity of the network). These type of decisions are of course slower. We can observe indeed that the fast part of the distribution comprises RT values that depend on the duration of the stimulus: in the upper row of fig.4.12 the fast part includes RTs of up to about 0.2 s, while in the bottom row the fast group includes RTs of up to about 0.6 s. In the bottom right panel the input due to the stimulus is always enough strong to drive the system to a decision in less than 0.5 s of stimulation. This phenomenon can be explained by the presence of the third "mixed" attractor. This stable point attracts the dynamics of the network even in the stimulus driven regime when the stimulus input is not too strong. When the system enters into the basin of attraction of the "mixed" state it will remain there until the stimulus estinguishes destabilizing the attractive point; in this case only after stimulus offset the network is eventually going to enter into
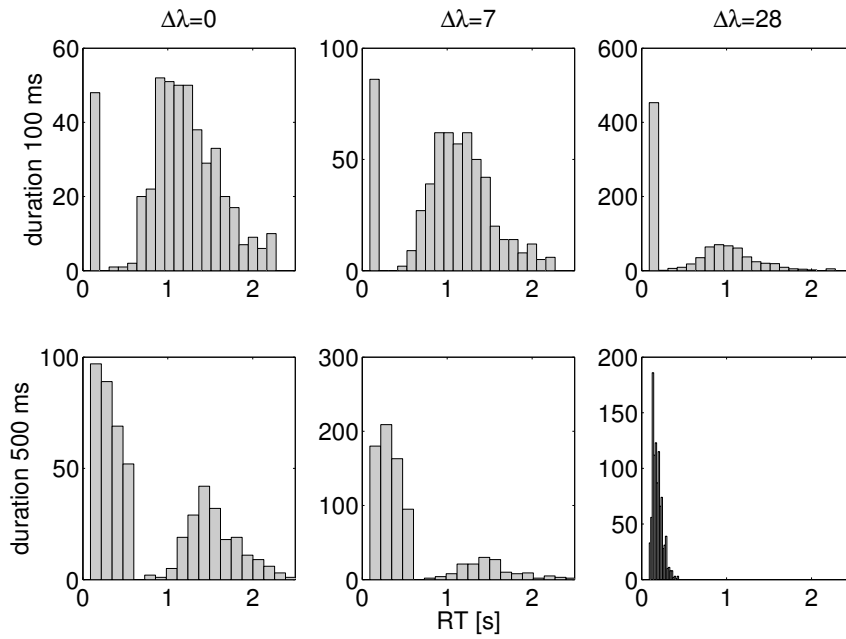
Figure 4.12: Reaction times histograms for three different values of $\Delta\lambda$ and two different durations. In these trials the sure target was not presented.

the basin of attraction of one decision. As a confirmation we observe that the RTs distributions are not bimodal if we set the parameter in order to work in a bistable regime (e.g. $\lambda = 15$Hz, figs. 4.20 and 4.20).

We now show what happens when the "sure" target is presented. As can be seen in fig. 4.13, the slow part of the distribution has always a smaller mass compared to fig.4.12 (when the "sure" target was not presented).

This difference arises because the trials in which the network doesn't reach a decision during the stimulus period and would thereby contribute to the slow part of the distribution, have a higher probability to end up in a "sure" decision.
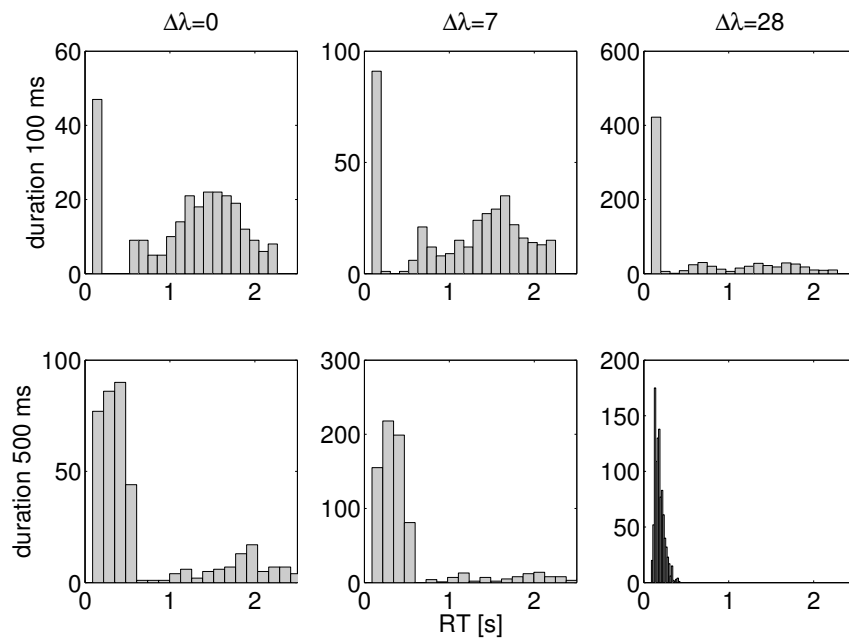
Figure 4.13: Reaction times histograms for three different values of $\Delta\lambda$ and two different durations. In these trials the sure target was presented.

4.2.6 Model Predicts Less Risky Behavior Approaching Third Bifurcation

Here we show the behavior of the model for different values of parameter λ, the common input to the decision pools. This parameter can embrace many complex phenomena (e.g. stimulus contrast, attention, urgency, etc.), that we didn't modeled explicitly since they seems to take place outside of the decision-making network. One of the well known effects of manipulating this parameter is the change in performance and speed of the network (Roxin and Ledberg, 2008a; Deco and Rolls, 2006). Then it is belived that the subjects have the ability to change, to some extent, their λ in order, for example, to satisfy different sets of experimental instructions ("fast" *versus* "accurate" trials) or for their own convenience.

Our psychophysical measurements of network activity (shown above in fig. 4.3) seem to indicate that the subjects in the "sure target"

experiment by Kiani and Shadlen (2009) set their working point in the multistable region near the third bifurcation, since we found the best fit of the model to experimental data in this region. Given that decision-making models usually work in the bistable regime (Wang, 2002b; Albantakis and Deco, 2009b; Marti et al., 2006) and the "mixed" attractor is not an useful representation for decision states, one could ask why the subjects should work in the multistable region. In order to address this issue we analyze in fig. 4.14 the overall $P(correct)$, the $P(S)$ and the consequent reward amount for each point the space of the input (λ).
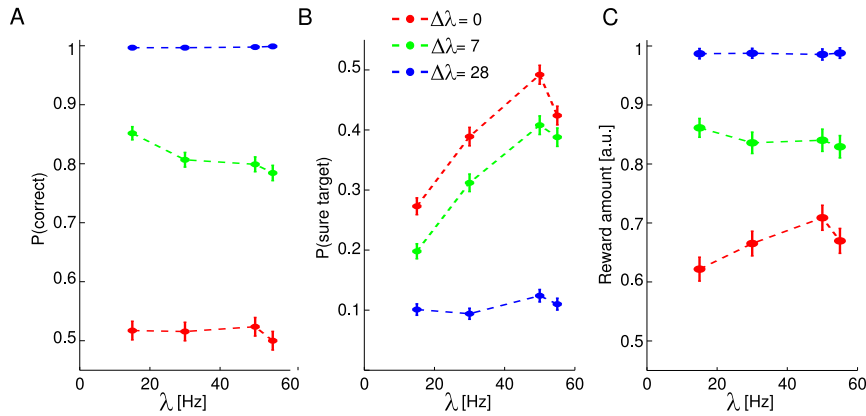
Figure 4.14: Fraction of correct responses, "sure" responces and reward rate.

In contrast to preavious analysis summarized in fig. 4.3, here we keep the values of $\Delta\lambda$ fixed when varying the value of λ, since we want to simulate the scenario of a subject moving its working point (due to experimental manipulation, like time pressure, stress on accuracy, stimulus visibility, etc. or its own will). We reasoned that in this scenario the $\Delta\lambda$ shouldn't be adjusted for each value of λ since, once the subject has learned the task, $\Delta\lambda$ should only depend on stimulus coherence.

For $\Delta\lambda = 0$ no significant difference appear as a function of λ. The same is true for the highest value of $\Delta\lambda$. This is comprehensible since no improvement in $P(correct)$ is possible when there is no evidence for the decision or when $P(correct)$ already plateaued. For intermediate

values of $\Delta\lambda$ the probability of correct choices decreases as a function of λ (see fig. 4.14A). Therefore it doesn't seem convenient for the subject to choose a working point near the third bifurcation.

But in the context of the present task the $P(correct)$ is not the only ingredient for an optimal behaviour.Indeed, in this task other valuable options are presented during the decision process. In this task, flexible decision-making is important in order to benefit from the "sure" option. As can be seen in fig. 4.14B the $P(S)$ increases as a function of the common input to decision pools and reaches a maximum at 50 Hz, in the proximity of the third bifurcation point. Near this point the stable mixed attractor holds the network longer in the undecided region, hence the decision is not yet well formed and the "sure" option can be taken into account.

Fig. 4.14C shows the resulting amount of reward from $P(correct)$ and $P(S)$. The amount of reward received is higher in the multistable regime and presents a maximum at 50 Hz.

Clearly the reward amount depends on the actual value assigned by the subject to each option. The more risk averse they are, the more they will prefere to choose a sure option at the expenses of performance. At this point we can infer that subjects in this experiment choose a working point near the pitchfork bifurcation because they were not willing to risk: they prefer to be flexible and make the most of the sure option. This result turns into the testable prediction that by changing the relative value of sure and high stakes options (e.g. promoting a risk seeking behaviour) the optimal working point would lay near the second bifurcation and the $P(S)$ would change consistently.

4.3 Discussion

4.3.1 Multiple Choice or Confidence Judgment?

We presented a biologically detailed network model that can account for behavioral and neurophysiological measurements related to decision confidence. Decision confidence is usually understood as a monitoring process (**?**) estimating the probability that the outcome of the decision

is correct (Kepecs and Mainen, 2012). This estimate can then be used to produce a subsequent decision (Kepecs et al., 2008; Insabato et al., 2010; Pleskac and Busemeyer, 2010) about the fist one (e.g. abort the action, change strategy, etc.). Different experimental paradigms have been used to measure the confidence in a decision. Besides the classical direct confidence rating, already used in the dawning of experimental psychology (Pierce and Jastrow, 1884) and possible only with humans, all other tasks can be classified in two types: uncertain-option tasks and post-decision wagering tasks (for a good review see (Kepecs and Mainen, 2012)). While post-decision wagering paradigms require an overt decision and a subsequent wagering of some form about this decision, in the uncertain-option paradigm on each trial either the perceptual decision or the confidence based decision are made explicit. The task used for our model (derived from Kiani and Shadlen (2009)), falls in the category of "uncertain-option". When the stimulus is presented the perceptual decision about the direction of dots motion needs to be computed but the response is delayed and the decision must be held in memory; then, before making an overt response, the subjectively estimated outcome of the decision has to be compared with the value of a "sure" option. Given the peculiar structure of this task and the associated neural recordings shown by Kiani and Shadlen (2009) a two stage model like the one proposed by Insabato et al. (2010) wouldn't be well suited. We therefore interpret the task as a delayed form of a triple choice decision making where the "evidence" for the third choice is not perceptual but rather based on the value of the third option. This interpratation is actually a common criticism moved against uncertain-option paradigm (Shields et al., 1997; Smith et al., 2006), which states that this task can be solved just by learning the appropriate stimulus responce association. Nonetheless we think that this task really reflects the confidence associated with the perceptual decision. Our opinion is supported by the fact that, when they are free to choose between the more valuable (but more risky) perceptual decision and the "sure" option, both monkeys and the model reach higher $P(correct)$ respect to forced-choice perceptual judgment trials. This means that they can access information about the expected outcome of the judgment and hence opting for the "sure" target denote low confidence in the decision (as already noted by Hampton (2001c)). Nonetheless an alternative explanation of the increased performances is possible,

which weaken the link of this behavior with confidence. Indeed fluctuations in the general vigilance state of the subject could also produce higher $P(correct)$ in free choice trials (since subjects would accept the perceptual task only when they have high vigilance). Although this would still imply a metacognitive process it would be different from a confidence judgment. In this regard our model produces an interesting prediction that can disantangle the question. The probability of choosing the "sure" target presents a mirror modulation when separated in early correct and error trials (fig. 4.11). This pattern is a hallmark of confidence (Kepecs et al., 2008) and can not be explained by fluctuations of vigilance, since vigilance is not expected to vary depending on the outcome of the decision. This prediction could be nicely tested (both in human and animal behavioural experiments) in variations of this task, where the response should be given with slow reaching motion towards the target. In a similar setup it could be possible, with the use of a traking device (eg. a vBOT manipulandum for primates (Howard et al., 2009; Resulaj et al., 2009) or a choice ball for rats (Sanders and Kepecs, 2012)), to track the initial decision of the subject before she choose the "sure" option.

4.3.2   The Timing of the Third Input

Given the nature of the mechanism implemented in our model it is worth asking if the timing of the third target has some impact on the decision. In the experiment Kiani and Shadlen (2009) varies the timing of presentation of the "sure" target from 500 ms to 750 ms after the extinction of the stimulus. This variation is introduced to avoid the predictability of the target but authors don't comment about any effect of this timing. For this reason we used a fixed delay between stimulus extinction and "sure" target presentation. Nevertheless we note that the mechanism we are proposing relies essentially on a dynamic effect and therefore the timing of presentation of the third target could influence the results. We reason that, for example, when the memory of the decision (during the delay period) is weak, it is going to be weaker and weaker as time elapses; in this condition a longer delay would produce more "sure" choices than a shorter one. The memory of the system is represented by the decision attractor that are stable

without stimulus input to the system. In this condition however the spontaneous state is also stable and, if the network didn't fall into the basin of attraction of one decision, it could wander around the diagonal and eventually be attracted by the spontaneous resting state. In this case it is more probable that a "sure" decision is going to be made as shown in fig.4.6. Anyway if Kiani and Shadlen (2009) didn't find an effect of the delay duration on the probability of "sure" responses, this could be due to the fact that the range of delay durations was too small. Nontheless very long delays could imply different neural structures and mechanisms respect to the ones proposed here and could supply distinct results. This and other issues related with the timing of the third target could be addressed in a future work.

4.3.3 Differences between Confidence Models

Although the race model could seem a more natural option to account for confidence processing (since it was already used to this aim (Vickers, 1979b)), Kiani and Shadlen (2009) used a one-dimensional diffution model to explain their data. They modified the decision rule of the standard DDM assigning "sure" choices if the accumulator was in a region near the mid-point of the diffusion space and one or the other perceptual choices otherwise (depending on the sign of the decision variable). This choice mechanism could be used because the system had to wait until the end of the trial and no self stopping mechanism was required. The model reproduces quite well the psychophysics experimental data. Nonetheless this model couldn't account for confidence processing in a RT task. Indeed in a RT task the model needs a threshold for stopping the accumulation but this would imply the same confidence judgment in all trials (unless a time-varying boundary is introduced (Drugowitsch et al., 2012)). Moreover this model is only able to reproduce neurophysiological findings during the stimulus period (since the DDM has no working memory). We also note that the DDM of Kiani and Shadlen (2009) does not provide novel testable predictions, which are foundamental to confirm or falsify a model.

On one hand diffution models have the appealing of being simple and producing neat explanation of observables; on the other they are only

phenomenological models and have the drawback of incorporating too few details about the biology of the systems modelled. Therefore chasing the goal of a mechanistic explanation of confidence processing in an uncertain-option task we propose the biophysically detailed network model presented above. Our model produces novel predictions that can be tested in new experiments: 1) an opposite modulation of the probability of the "sure" option in correct and error trials, when plotted against the easiness of the trial; 2) a bimodal distribution of RTs (that could be tested in a variation of the task as explained above); 3) a different modulation of the probability of the "sure" choices (and overall a smaller amount) when the input to the network is smaller; this can correspond to the condition of a more risk-seeking behaviour or less time pressure in a RT version of the task.

We also want to distinguish the present model from the one proposed in the second chapter (Insabato et al., 2010). Both are ANN models of confidence processes but they differ in the structure of the network and the underlying computations. The network proposed in chap. 2 has a two layer structure. The perceptual decision is effected in a module that implement a biased competition mechanism (Wang, 2002b) between two decision selective neuronal pools. When the discriminability of the stimulus is increased the decision attractor of the correct choice moves towards higher firing-rate of the winning pool. On the contrary the decision attractor associated with error choices moves towards lower firing-rates of the winning pool. Therefore the sum of firing-rates of decision pools in this module produces the characteristic x-shaped pattern observed in confidence studies. The confidence judgment in this network is effected in another module, which also implement a biased competition mechanism. But this module does not receive the stimulus input, rather the "high confidence" pool gets as input the activity of decision pools in the percetual module, while the "low confidence" pool gets a reference signal with a fixed mean.

The model presented in this chapter instead has just one layer because it absorbes confidence judgment process into the same perceptual decision module (this reflects the structure of the behavioral task, where there is only one final overt binary response, i.e. either "high confidence - left/right" or "low confidence"). In this network first the pools associated with the perceptual choice compete in order to reach a de-

cision about the motion direction of RDM stimulus and then, when the "sure" target is turned on, a new competition is triggered between the pool associated with the "sure" target and the other two decision pools.

Surely there are certain similarities between the models. 1) The perceptual decision is taken using a biased competition mechanism. 2) The pool representing a decision of low confidence (the S pool in the model of this chapter) only receives a mostly value encoding signal with a fixed mean. 3) The high confidence decision is driven by the activity of decision pools. Nonetheless the two models are different under other aspects. 1) The confidence judgment is made in a separate network with its own dynamics for the model in chap. 2. 2) The low confidence pool competes only indirectly with the perceptual decision pools in the model of chap. 2. 3) In the model of this chapter the confidence is encoded in an non linear distance between the firing-rates of decision pools, while we claim in chap. 2 that the sum of firing-rates is the vehicle of confidence.

We note that this discrepacy of the confidence encoding in the two models could be possibily reconciliated by using an n-norm for the two-layer model. Indeed, while the second network is receiving the sum of the activity of perceptual decision pools (a 1-norm) the processing in the confidence network could produce a confidence encoding compatible with a higher order norm. The maximum norm (for $n = \infty$) is equal to $\max\{\nu_L, \nu_R\}$ and is therefore equivalent to the distance function used in this chapeter. Therefore in the case of a high value of n the sum and the difference of the activity of decision pools are approximately the same.

As a last remark, although the model of chap. 2 could in principle account for the data of Kiani and Shadlen (2009), it seems a more natural decision to use a multiple choices decision-making network for the present task. Indeed the structure of the task is very similar to that of a three alternative choice task and what really changes is the timing of apparence of the targets and the associated value and probability of reward. Following Occam's razor, we used here a more minimalistic approach and proposed that the confidence in a decision can be extracted and used in the same simple multiple choices decision

making network exploiting a multistable regime.

4.4 Methods

4.4.1 Model details and mean-field reduction

We run simulations of 1000 trials for each stimulus condition (duration and coherence or $\Delta\lambda$). The stimulation protocol of the selective pools R and L is the following. First during 500 ms the pools only receives the background activity to assure that the network is in resting state. Then the target input is delivered similarly to other studies (e.g. Albantakis and Deco (2009b); Furman and Wang (2008)). The target input is modeled as the concatenation of two exponential: in the first 400 ms it was set to $200 + 100exp(-t/100)$ Hz, after that during 100 ms it is supposed to decay according to $200exp(-t/15)$ Hz. Then the stimulus input is turned on: pool L receives $\lambda + \Delta\lambda$ while pool R receives $\lambda - \Delta\lambda$. After the stimulus estiguishes (according to ist duration, which was either 100, 200, 300 or 500 ms) λ is set to 0 until the end of the trial. At the end of the trial a saccade related signal of 80 Hz is delivered for 100 ms. The pool S receives only the background activity until the "sure" target is turned on (500 ms after the end of the RDM stimulus). The input due to the "sure" target is modeled as the target input to pools R and L ($40 + 200exp(-t/100)$ Hz) but it reached a higher steady value of $\lambda_{sure} = 40$ Hz that encodes the value of the "sure" option. The stimulation protocol is show in fig.4.1.

Detailed description of the neuron and synapse model is given in the appendix A and in tabs. 4.1 and 4.2.

We use a mean-field reduction (Brunel and Wang, 2001b) in order to study the space of the principal parameters of the network.

The mean-field approximation allows reduction of the number of dynamical variables, by describing the average firing rate of each neuronal pool in the limit of an infinitely large number of neurons. The network dynamics Although our network has three decision pools, the stimulation protocol is asynchronous: in the first phase of the stimulation the pool S receives only the background activity and its firing rate remains

A	Model Summary
Populations	five
Connectivity	full, no synaptic delay
Neuron model	Leaky Integrate-and-Fire, fixed threshold, fixed refractory time
Synapse model	Instantaneous jump and exponential decay for AMPA and GABA and exponential jump and decay for NMDA receptors
Plasticity	-
Input	Independent fixed-rate poisson spike trains to all neurons
Measurements	Spike activity

B	Populations		
Total number of neurons	$N = 1000$		
Excitatory neurons in each module		$N_E = 0.8 \cdot N$	
Inhibitory neurons in each module		$N_I = 0.2 \cdot N$	
Name	**Size**	**Name**	**Size**
L (decision "left")	$N_L = f \cdot N_E$	Nonspecific	$N_E - 3N_A$
R (decision "right")	$N_R = f \cdot N_E$	Inhibitory	$0.2 \cdot N$
S (decision "sure target")	$N_S = f \cdot N_E$		

Table 4.1: Model summary. Characteristics of the model and populations details. Parameters values are given in Tab. 4.3

below 2 Hz. Therefore in this phase the system behaves like a binary decision-making network. The network activity could converge to one of four attractors: a spontaneous state, where the selective pools have low activity; two selective states with one selective pool firing at a high level and the other inhibited; and a mixed state with both pools highly active. With some parameter values the network could sustain just one stable state, or could show bistable, behaviour, or could show multistable behaviour in which the spontaneous state or the mixed state and each of the two decision states are all possible stable states when the inputs to the network are being applied (Brunel and Wang, 2001b). Details of the mean-field analysis are given in the appendix B.

We also calculate the boundary of the attraction basins of fixed points (shown in fig.4.4). In order to find the boundaries we set the initial values of the firing rate of pools L and R to different points in the plane

C	Neuron and Synapse Model
Type	Leaky integrate-and-fire, conductance-based synapses
Subthreshold dynamics	$C_m \dot{V}(t) = -g_L(V(t) - V_L) - I_{AMPA,ext}(t) - I_{AMPA,rec}(t) - I_{NMDA}(t) - I_{GABA}(t)$
Synaptic currents	$I_{AMPA,ext}(t) = g_{AMPA,ext}(V(t) - V_E)\sum_{j=1}^{N_{ext}} s_j^{AMPA,ext}(t)$ $I_{AMPA,rec}(t) = g_{AMPA,rec}(V(t) - V_E)\sum_{j=1}^{N_E} w_j s_j^{AMPA,rec}(t)$ $I_{NMDA}(t) = \frac{g_{NMDA}(V(t)-V_E)}{1+[Mg^{2+}]exp(-0.062V(t))/3.57} \times \sum_{j=1}^{N_E} w_j s_j^{NMDA}(t)$ $I_{GABA}(t) = g_{GABA}(V(t) - V_I)\sum_{j=1}^{N_I} w_j s_j^{GABA}(t)$
Fraction of open channels	$\frac{ds_j^{AMPA,ext}(t)}{dt} = -\frac{s_j^{AMPA,ext}(t)}{\tau_{AMPA}} + \sum_k \delta(t - t_j^k)$ $\frac{ds_j^{AMPA,rec}(t)}{dt} = -\frac{s_j^{AMPA,rec}(t)}{\tau_{AMPA}} + \sum_k \delta(t - t_j^k)$ $\frac{ds_j^{NMDA}(t)}{dt} = -\frac{s_j^{NMDA}(t)}{\tau_{NMDA,decay}} + \alpha x_j(t)(1 - s_j^{NMDA}(t))$ $\frac{dx_j^{NMDA}(t)}{dt} = -\frac{x_j^{NMDA}(t)}{\tau_{NMDA,rise}} + \sum_k \delta(t - t_j^k)$ $\frac{ds_j^{GABA}(t)}{dt} = -\frac{s_j^{GABA}(t)}{\tau_{GABA}} + \sum_k \delta(t - t_j^k)$
Spiking	if $V(t) \geq V_\theta \wedge t > t^\star + \tau_{rp}$ 1. $t^\star = t$ 2. emit spike at time $t^\star$ 3. $V(t) = V_{reset}$

D	Input
Type	**Description**
Poisson generator	Fixed rate, N_{ext} poisson generators per neuron, each one projects to one neuron

Table 4.2: Model summary. Neuron model and input layer description. Parameters values are given in Tab. 4.3

Parameter	Value	Parameter	Value
C_m (excitatory)	0.5 nF	V_E	0 mV
C_m (inhibitory)	0.2 nF	V_I	-70 mV
f	0.20	V_L	-70 mV
$g_{AMPA,ext}$ (excitatory)	2.08 nS	V_{reset}	-55 mV
$g_{AMPA,ext}$ (inhibitory)	1.62 nS	V_θ	-50 mV
$g_{AMPA,rec}$ (excitatory)	0.104 nS	w_+	1.8
$g_{AMPA,rec}$ (inhibitory)	0.081 nS	w_-	0.878
g_{GABA} (excitatory)	1.287 nS	α	0.5 ms^{-1}
g_{GABA} (inhibitory)	1.002 nS	λ_{sure}	40 Hz
g_{NMDA} (excitatory)	0.327 nS	λ_{ext}	2.4 kHz
λ	$\{15,30,50,55\}$ Hz	$\Delta\lambda$	$[0,28]$ Hz
g_{NMDA} (inhibitory)	0.258 nS	τ_{AMPA}	2 ms
N_E	800	τ_{GABA}	10 ms
N_I	200	$\tau_{NMDA,decay}$	100 ms
N_{ext}	800	$\tau_{NMDA,rise}$	2 ms

Table 4.3: Parameters used in the simulations.

ν_L, ν_R and record which attractor the network falls in; we change the initial values with a bisection algorithm until finding the boundary with a precision of 0.1 Hz.

4.4.2 Probability of "sure" target selection: $P(S|\nu_L, \nu_R)$

We calculate the probability of choosing the "sure" option conditioned on the value of firing rate of decision populations R and L. In order to calculate this probability we devided the plane ν_L, ν_R in square bins of side 1 Hz. We calculated for each trial the time averaged firing rate of decision pools in the 50 ms time window preceding the "sure" target onset. Then each bin in the plane is associated with those trials presenting the average firing rate before "sure" target onset determined by the position of the bin. Finally we counted for each bin how many trials ended up in a "sure" decision and divided by the total number of trials associated with that bin. We did this analysis separately for each stimulus duration collapsing together all coherences ($\Delta\lambda$) and

also for each coherence collapsing together all durations. The results were very similar across different conditions therefore we decided to pool together all trials from all stimulus conditions in order to show cleaner results.

4.4.3   Undecided time

We show that the system spends more time in the undecided region when stimulus conditions (duration and coherence) determine less confidence on average. We defined the undecided region to be the region of the plane ν_L, ν_R enclosed in the convex hull of points corresponding to $P(s|\nu_L, \nu_R) > 0.6$. Of course the value chosen (0.6) is arbitrary and other values give similar results, provided that they are not too high or too low. We then calculate the time that the system spends in the undecided region from stimulus on until the "sure" target is turned on.

4.4.4   Reaction time

We calculate the decision time of the network as the time when the firing rate of a decision pool crosses a threshold set at 25 Hz. In addition the firing rate need to be above the threshold in the subsequent 50 ms. Of course this values are not meaningful and other similar values give the same qualitative results. The reaction time is the sum of the decision time and a non-decision time that is usually used to account for sensory processing and motor planning. We set the non-decision time to 80 ms but our results do not depend on this value.

4.4.5   Reward amount

We calculate the reward amount received by the subjects as a function of the probability of correct response and the probability of "sure" target selection. We arbitrarily consider the value of a correct response to be 1 and that of an error to be 0. The value of the "sure" target has been set to 0.8, since in the experimental setup the subjects received for a "sure" response approximately 80% of the reward associated to the

high stakes targets. However we note that this value neither is crucial for our results nor need to be fixed based solely on the experimental setup. Indeeed, while the value attributed by the subject to the "sure" option will reflect the objective amount of reward associated, it won't probably be the same since it will include subjective preferences like risk aversion or risk attraction. We calculated the total reward in a block of trials as the sum of the reward received in each trial according the choice made (error: 0; correct: 1; sure: 0.8). Reward functions can be complicated and include many variables like reaction times, penalty time associated to errors, etc. (Drugowitsch et al., 2012). We decided to use a very simple function since each of these variables need to be weighted based on the subjective value assigned to it, as for the value of the "sure" option, and we don't have information about these subjective values.

4.A   Supplementary figures
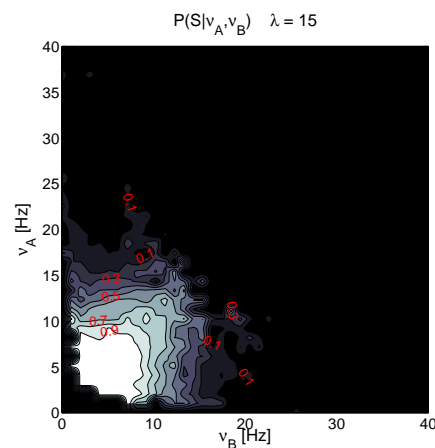
Figure 4.15: Conditional probability of choosing the "sure" option given the state of decision neurons for $\lambda = 15$. The color convention is the same of fig.4.6.
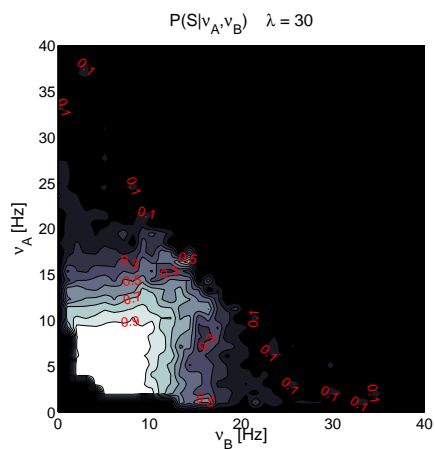
Figure 4.16: Conditional probability of choosing the "sure" option given the state of decision neurons for $\lambda = 30$. The color convention is the same of fig.4.6.
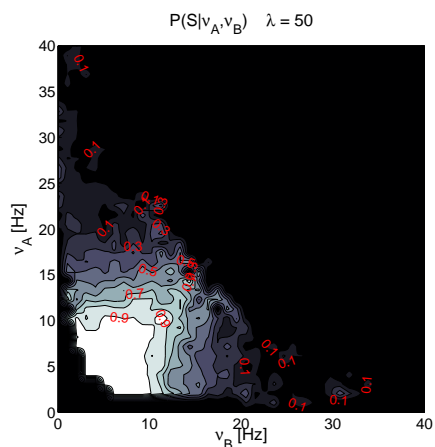
Figure 4.17: Conditional probability of choosing the "sure" option given the state of decision neurons for $\lambda = 50$. The color convention is the same of fig.4.6.
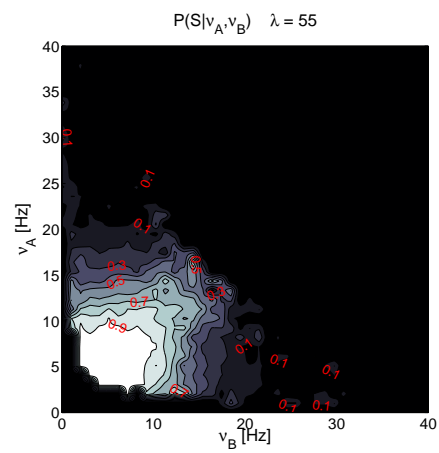
Figure 4.18: Conditional probability of choosing the "sure" option given the state of decision neurons for $\lambda = 55$. The color convention is the same of fig.4.6.
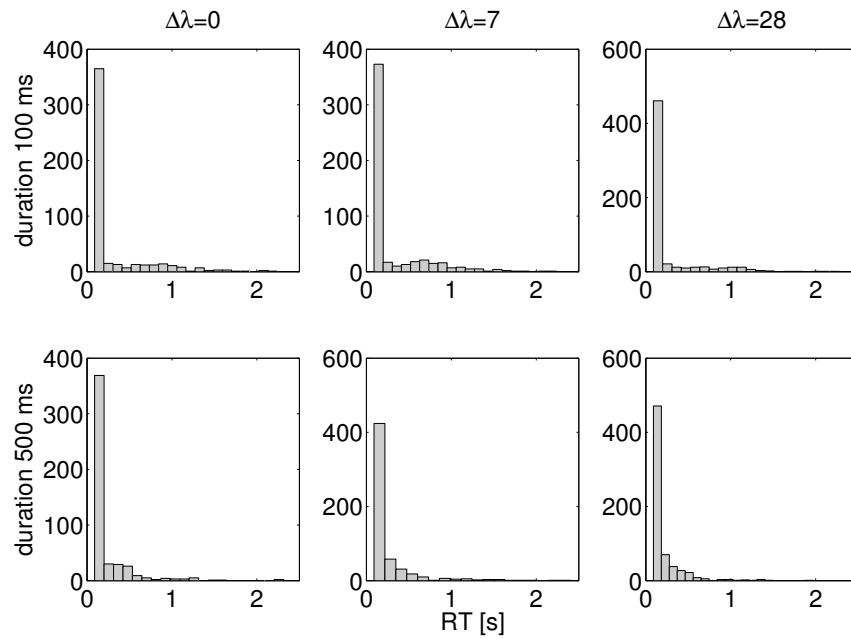
Figure 4.19: Reaction times histograms for three different values of $\Delta\lambda$ and two different durations. The average common stimulus intensity was $\lambda = 15$ Hz. In these trials the sure target was not presented.
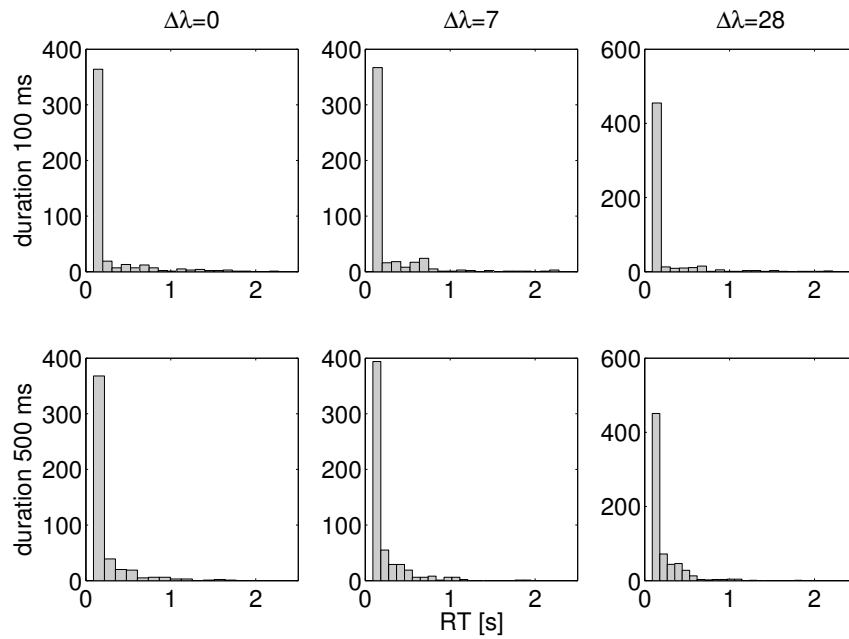
Figure 4.20: Reaction times histograms for three different values of $\Delta\lambda$ and two different durations. The average common stimulus intensity was $\lambda = 15$ Hz. In these trials the sure target was presented.

Coda

> ... wine is good, and confidence is good; but can wine or confidence percolate down through all the stony strata of hard considerations, and drop warmly and ruddily into the cold cave of truth?
>
> H. Melville, The confidence-man

In this chapter we will briefly resume the results presented in previous chapters. We will then highlight several aspects that are relevant for the study of decision confidence and need to be discussed.

5.1 A Neurocomputational Framework for Decision Confidence Studies

Kepecs and Mainen (2012) already argued that we need a computational framework for understanding decision confidence. Surely we do not pretend to have built this framework but at least we hope to have given in the previous chapters some building blocks for constructing it.

In chap. 2 we proposed a model that accounts for neural processing of decision confidence in post-decision wagering experiments. We based our model on the findings of Kepecs et al. (2008). They showed that neurons in rat OFC present the characteristic X shaped modulation of confidence measures. They also demonstrate that rats are able to make confidence judgments in a post-decision wagering experiment. We modeled the neural system involved in this task as a two layers network. The network is composed of a decision-making module and a confidence module. The decision-making module receives the stimulus related input and makes the perceptual choice. The confidence module receives the activity of decision-making neurons and an external signal, related to the value of options, and takes a second decision: to stay and wait for reward or to leave and start a new trial. Our model is able to reproduce both the neural activity and the behavior of rats. Moreover the model makes the prediction that the OFC modulation is the result of averaging together trials with high confidence and trials with low confidence.

Building on these ideas in chap. 3 we analyzed neurons from monkeys PMv recorded during a perceptual decision-making task. Our aim was to study the activity of real neurons in order to support or falsify the model of chap. 2 and to shed light on decision confidence encoding. Specifically we wanted to understand whether neurons encode confidence in a continuous or discrete way. Indeed, while the X shaped pattern is well established and has been found in rats OFC activity (Kepecs et al., 2008), it is a modulation of average confidence signals in block of trials distinguished by outcome and difficulty. We wanted to ask what happens in single trials. To this aim we first identified neurons presenting the X shaped modulation by means of linear regression analysis. We reasoned that the modulation of average activity as a function of the difficulty of the task could arise from at least three different mechanisms, each of them corresponding to a different confidence encoding in single trials: a switch timing encoding, a firing-rate encoding and a binary encoding. The switch timing code hypothesis states that when neurons switch from a resting state to an active state that produces the confidence signal, the modulation of activity does not depend on the firing-rate of the active state but on the timing of the switch. On the contrary the firing-rate code imply that

there is a correlation between the single trial firing-rate after the switch
and the difficulty. The binary code corresponds to the mechanism pre-
dicted by the model of chap. 2. On single trials neurons encode either
high or low confidence and the proportion of high (or low) confidence
trials correlates with the difficulty. These three mechanisms are not
exclusive. Indeed we found different neurons implementing each of
the three mechanisms. This rich representation of decision confidence
in the brain can be useful since continuous codes (switch timing and
firing-rate) can well represent confidence but a discretization stage is
needed when a categorization of confidence is required (e.g. a stay or
leave choice). We note therefore that the binary neurons we found in
monkey PMv support the model of chap. 2. On the other hand the
neurons using continuous codes, firing-rate and switch timing code,
are not predicted by the model. Nonetheless we think that a simple
modification of the network could account for that neurons. Indeed in
the model the decision confidence is implicitly encoded in the sum of
the activity of decision-making neurons. An additional stage explicitly
encoding confidence could be obtained by creating synaptic projections
of decision-making neurons to a new pool of neurons that would then
project to the confidence module.

While in chap. 2 and in chap. 3 we dealt with the post-decision wa-
gering task in chap. 4 we considered the unsure option task. In this
chapter we used the findings of Kiani and Shadlen (2009) as constrains
for the model. Kiani and Shadlen (2009), using the unsure task and
the random dot motion stimulus, showed that LIP neurons in primate
brain present an intermediate level of firing-rate associated with un-
sure trials compared with sure trials. We built a model with a multiple
choice decision-making architecture that can account for their findings.
In our model two decision selective pools encode the decision about the
direction of motion while a third pool encode the "sure" choice. When
the "sure" target is not presented the third pool remains silently in
resting state and do not affect the decision process. When the "sure"
target is presented a competition is triggered between the "sure" pool
and the decision-making pools. We found that the confidence measure,
i.e. the probability of choosing the "sure" target, has a non-linear rela-
tionship with the activity of decision pools. Indeed this probability is
related to a non-linear measure of the distance between the activity of

decision pools. This differs from the classical account of confidence in race models (Vickers, 1979b) where the 1-norm distance (absolute difference) between accumulators is used. We also found that the model is able to reproduce the experimental findings in a multistable regime in the proximity to a pitchfork bifurcation, where there is a stable fixed point corresponding to an undecided state with both decision pools firing at a high rate. This model produces interesting predictions that can be tested in new experiments. 1) A mirror modulation of the confidence measure is predicted in early correct compared to early error trials. This prediction, if confirmed, would remove the doubts raised about this task not reflecting confidence but other metacognitive processes. 2) Reaction times have a bimodal distribution since in "unsure" trials the system is kept next to the undecided fixed point and can reach a decision attractor only when the undecided attractor looses stability. 3) Some manipulations of the task, e.g. induction of risk-seeking behavior, should push the subject to move away from pitchfork bifurcation (and possibly from the multistable region) and produce different behavior and neural activity as predicted by the model.

5.2   *Ad Ventura*

In this section we want to summarize the possible future development of the research about decision confidence.

The first natural follow-up study would be the development of an experimental setup to support or falsify the model of chap. 4. The model predicts a different behavior if the system is pushed to work in a bistable regime. This could be done by diminishing the desirability of the "sure" option. Indeed in this case the subject should try to work in a regime where the probability of correct responses is higher (see fig. 4.14). In general any manipulation that fuel a risk-seeking behavior should bring the system to the bistable regime.

The other two predictions of the model are probably more complicated to reproduce in an experimental setup since they involve the measurement of early direction of the decision process and reaction time. However we think it would be feasible if the subjects are re-

quired to use slow reaching movements to select the targets. This type of movements has been already used in decision-making both with rats (Sanders and Kepecs, 2012) and primates (Howard et al., 2009; Resulaj et al., 2009). If the "sure" target appears when the selection movement has already been started but not completed (at least in most trials) the initial decision and the reaction time could be recorded.

Another interesting extension of the research presented here would be a computational account of confidence judgments beyond binary classification. Both the model of chap. 2 and that of chap. 4 only give binary judgments of decision confidence, because they are meant to account for tasks where only two categories of confidence were used. However with human subjects it is usual to adopt confidence scales with more than two options (refer to sec. 1.2.3 for examples). Although multiple choices mechanisms have been studied in the context of ANN (Albantakis and Deco, 2009b), it is not trivial to extend our models to multiple options. Moreover the extreme case of a confidence judgment on a continuous scale (Graziano and Sigman, 2009; Barttfeld et al., 2013) places an even more complex challenge.

In addition an intriguing idea comes from the study of Zylberberg et al. (2012). They studied the temporal evolution of decision confidence in simple perceptual decision tasks by constructing the confidence kernels (similarly to the well known decision kernels (Kiani et al., 2008)). Their results can be used as constrain for the models. Indeed they fitted three different phenomenological models to the decision kernels of subjects: a race model, a one-dimensional diffusion model, a partially correlated race model and the "Vicker's model" (inspired by Vickers (1979b)). Only the partially correlated model could reproduce at least qualitatively the confidence kernels. The study of our model of chap. 2 in the light of these findings could be a good starting point, since it could highlight what the meaningful parameters are that produce this type of temporal evolution of decision confidence.

5.3 One Model?

In the field of decision-making studies in neuroscience there is a never ending dispute between the fans of the two major theoretical framework, the DDM and ANN. Several models have been used to account for decision confidence processing both in the ANN and DDM framework. We want to reflect here about this dichotomy between ANN and DDM taking advantage of their results about decision confidence.

A first question that we want to ask is: are all these models different? It could seem silly to ask such a question but actually even the two frameworks are not so distinct. Indeed Roxin and Ledberg (2008a) demonstrated that a decision-making ANN model can be reduced to a one-dimensional non-linear diffusion equation for a specific set of parameters (near a bifurcation). This result indicate that at least for a given dynamical regime the ANN is reducible to a diffusion process (although with a non-linear term).

Given this result it seems not so trivial to answer the question.

Probably the very fundamental difference between the two frameworks is the level of description of the phenomena. While the DDM describes the decision as the diffusion of a particle the ANN describes it as the result of the interaction of thousands of neurons in the brain. Although this difference is only semantical it is important to reflect on what it means to use one model instead of another. When we use an ANN to account for decision-making, we are implicitly implying that exist a network of neurons in the brain equal (or very similar) to that of the model and that the activity of that network is the cause of the behavior. When we use a DDM, are we suggesting that a particle diffusing in the brain of the subject is the cause of the choice? Or do we believe that the DDM is a useful formalism for describing the process and that some neural machinery in the brain implements it? The second interpretation seems more sound to us. In this case the ANN could be seen as a useful formalism to understand how a neural mechanism can implement a process similar to the diffusion of a particle.

The answer to the question is therefore not a simple one. We don't want to claim that DDM and ANN are the same but we want to focus

the attention on the fact that they should not be always considered as opposite.

The most meaningful way of comparing models is, in our opinion, the Popperian falsification protocol (as also advocated in a more general context by Sober (1994)), in which the predictions of each model are tested by means of new experiments.

5.3.1 Predictions

A good model should then formulate clear predictions that can be tested. The two layers model for confidence decisions of chap. 2 predicts a bimodal distribution of firing-rates in confidence neurons and we showed in chap. 3 that monkey PMv neurons confirm this prediction. Nonetheless we note that this confirmation should not reassure us that our model is correct for at least two reasons: 1. It needs further confirmation by other experimental studies; 2. We should try to produce more predictions to falsificate the model and push forward our understanding of decision confidence.

In the previous chapter we proposed an ANN model for confidence computations in LIP, as evidenced by the experimental data of Kiani and Shadlen (2009), that gives some testable predictions. It is interesting when two different models that account for the same set of data produce contrasting predictions. Kiani and Shadlen (2009) explain their data with a DDM as described in sec.1.2.3. Unfortunately they don't provide any novel testable prediction.

However interesting predictions about decision confidence comes also from the framework of DDM. For example Moreno-bote (2010) predicts that the relation between confidence level and RTs is not linear in time, like suggested by the usual fit $(a/t + b)$ found by many experiments (Henmon, 1911; Volkmann, 1934; Reed, 1951; Audley, 1964), but depend on the square root of time as $a/\sqrt{t} + b$. This prediction needs a carefully designed experiment to be tested since the difference between the two types of relationships is only visible with quite long RTs.

Another inspiring work is that of Zylberberg et al. (2012), where predictions are provided for four different versions of DDM (a one-dimensional diffusion model, a race model, a partially correlated model and the race model of Vickers (1979b)). They calculated the confidence kernel, by means of inverse psychophysical correlation analysis, for each model. This measure is similar to the more usual decision kernel (Kiani et al., 2008) and provides an estimate of the temporal relevance of the input/stimulus fluctuations on the confidence judgment. They then calculated the confidence kernels of subjects in a perceptual experiment and compared the predicted kernels with that obtained from experimental data.

None of the models was able to reproduce well the confidence kernels although the partially correlated model could adjust at least qualitatively to the experimental data. These results are particularly interesting since they allow to distinguish the different models on the basis of their predictions. Therefore it would be important to further understand the failure of the models in reproducing the confidence kernels and to test also the ANN models within this framework.

The available confidence models are also different regarding the proposed encoding of decision confidence. Vickers (1979b) proposed that the confidence could be encoded in the "balance of evidence", the difference between the accumulated totals of the two integrators of a race model. Moreno-bote (2010) showed that the confidence not only depends on the balance of evidence but also on the square root of the decision time. This encoding reduces to the sole time dependence in the case of the one-dimensional diffusion process since in this model the state of the second integrator is fully determined by the state of the first and in turn by the decision boundary. In chap. 2 we proposed that the decision confidence is encoded in the sum of the firing-rates of decision populations. This encoding arises from the architecture of the network, where the decision pools project to a population of neurons that encode the confidence in the decision. Therefore the confidence is encoded in the sum and not the difference between decision pools. In this model actually, when the decision has been taken, the firing-rate of the winning pool is the most informative since the loosing pool has very low activity due to the indirect inhibition of the other pool. Analyzing the model of chap. 4 we found indeed that the confidence is

encoded in that model in a distance measure between the activity of the decision pools. However this distance is not the simple balance of evidence and but a non-linear function of it: $\sqrt[n]{|\nu_L^n - \nu_R^n|}$, when n is large enough. This function is equal to $\max\{\nu_L, \nu_R^n\}$, for $n \to \infty$. This means that what really counts in this model for decision confidence evaluation is the activity of the winning pool. Only in a multistable regime, when the activity of the two decision pools are similar, their difference affects the confidence in the decision.

These distinct encoding of decision confidence don't provide testable predictions by themselves but can be used anyway to further understand the nature of confidence processing. Indeed when the confidence information encoded in a decision-making mechanism need to be used to produce a given behavior (e.g. a bet on the outcome of the decision), another mechanism has to read out that information. It would be interesting, for example, to study the possible mechanisms that can read out the decision time of a DDM.

5.4 The concept of decision confidence

In this dissertation we abstained from giving a direct definition of decision confidence. Surely we agree with Kepecs and Mainen (2012) that a formal computational foundation of the phenomenon of confidence is far more valuable than a semantic definition. Indeed especially the semantics surrounding those concepts that involve subjective experience (like consciousness or confidence) could be not adequate or sufficient to define the underlying phenomena (Metzinger, 2000). On the contrary a computational account aims at characterizing the system in terms of functions and computational units (Marr, 1982). From the account of decision confidence sketched in this work it seems that our concept of "confidence" in a decision should be probably revised or better detailed. We want to recall here the simplified schema of confidence processing outlined in fig.1.1. In that schema we separated the more abstract representation of the reliability of a decision, that we called "confidence" module, from the classification of this representation in discrete categories, that we called "confidence-based decision" module. This discretization stage is useful for translating the confi-

dence representation into a congruent action. The neurophysiological evidence provided in chap.3 about neurons continuously encoding confidence supports the idea of a module that represents the reliability of the decision process. On the other hand we provided both a computational account (in chap.2) and neurophysiological evidence (in chap.3) for a classification stage, where a second decision is taken based on confidence level in a first decision. Therefore, in the light of our results, it seems suitable to us to split the concept of decision confidence in at least two parts: The representation of the reliability of a decision and the classification of this sensation, be it a verbal rating or a post-decision wagering, we could also call it the confidence "judgment". This two entities correspond to the two stages of fig.1.1, the "confidence" module and the "confidence-based decision" module respectively. The decision confidence depends on the sensory input and the decision process and is connected to the subjective, phenomenal experience of deciding. Whereas the confidence judgment is of a discrete nature and is a function of the decision confidence, useful to communicate the confidence sensation to others.

Another aspect to be considered is the set of constraints imposed by the environment on the confidence judgment. For example it is plausible to assume different neural mechanisms if the possible confidence categories are only two versus five or, even more extremely, when the judgment has to be expressed on a continuous scale. As we use to talk about performance independently from task contingencies but we don't expect them to be produced by the same neural networks, the same could be true for confidence. Indeed we presented two different models in order to account for two types of confidence judgments: The post-decision wagering (chap.2) and the uncertain-option task (chap.4). The different tasks can therefore modify the simple scheme depicted in fig.1.1. In the first model there are two layers, one for taking the sensory decision, that corresponds to the "decision-making" module of fig.1.1, and another for taking a decision based on the confidence in the sensory decision. This last layer corresponds to the "confidence-based decision" module of fig.1.1. In this model the "confidence" module is not implemented explicitly but is integrated into the decision-making stage. Even more interesting in the model of chap.4 there is no separation between the decision-making, the representation of the reliability

of this decision and the confidence judgment. All three functions are integrated into the same neural machinery.

To conclude the idea, yet to be proved, that emerges from our study is that the concept of decision confidence should be separated in two concepts: the representation of the reliability of a decision, that is an internal feeling of confidence and the overt confidence judgment, that is manipulated by the way we express it.

Neuron and synapse model

At the neuronal level we used single compartment leaky integrate-and-fire (IF) neurons that incorporate biophysically realistic parameters Abeles (1991) without being computationally intractable Dayan and Abbott (2001); Brunel and Wang (2001b); **?**); **?**. The integrate-and-fire formulation is very important, for the spiking of the neurons in the network is close to random in its timing for a given mean firing rate (Poisson-like), and this introduces noise into the network which enables it to account for probabilistic decision-making, because the choices are influenced by the spiking-related coherent statistical fluctuations that are produced in a finite-size network.

The dynamics of the membrane potential in the subthreshold regime is determined by the membrane capacitance and a *leakage* term, following the equation:

$$C_m \dot{V}(t) = -g_L(V(t) - V_L) - I_{syn}(t), \qquad (A.1)$$

where C_m is the membrane capacitance, g_L is the membrane conductance that provides the leakage current, V_L is the resting potential, and I_{syn} is the total input synaptic current. When the depolarization of the membrane reaches a threshold value V_θ, the neuron fires a spike, its voltage is reset to a given value V_{reset}, and a refractory pe-

riod τ_{rp} follows during which the neuron is unable to fire another spike.

The total synaptic current is the sum of the external excitatory currents mediated by AMPA receptors ($I_{AMPA,ext}$), recurrent excitatory currents mediated by AMPA receptors ($I_{AMPA,rec}$) and NMDA receptors (I_{NMDA}), and inhibitory currents mediated by GABA receptors (I_{GABA}):

$$I_{syn}(t) = I_{AMPA,ext}(t) + I_{AMPA,rec}(t) + I_{NMDA}(t) + I_{GABA}(t). \quad \text{(A.2)}$$

Each current is defined by:

$$I_{AMPA,ext}(t) = g_{AMPA,ext}(V(t) - V_E) \sum_{j=1}^{N_{ext}} s_j^{AMPA,ext}(t) \quad \text{(A.3)}$$

$$I_{AMPA,rec}(t) = g_{AMPA,rec}(V(t) - V_E) \sum_{j=1}^{N_E} w_j s_j^{AMPA,rec}(t) \quad \text{(A.4)}$$

$$I_{NMDA}(t) = \frac{g_{NMDA}(V(t) - V_E)}{1 + [Mg^{2+}]exp(-0.062V(t))/3.57} \times \sum_{j=1}^{N_E} w_j s_j^{NMDA}(t)$$
$$\text{(A.5)}$$

$$I_{GABA}(t) = g_{GABA}(V(t) - V_I) \sum_{j=1}^{N_I} w_j s_j^{GABA}(t), \quad \text{(A.6)}$$

where $g_{AMPA,ext}$, $g_{AMPA,rec}$, g_{NMDA} and g_{GABA} are the receptor specific synaptic conductances. V_E and V_I are the reversal potentials respectively for the excitatory and inhibitory neurons. N_{ext}, N_E and N_I are respectively the number of external neurons residing in other cortical areas, of excitatory neurons and of inhibitory neurons. $s_j^{AMPA,ext}, s_j^{AMPA,rec}, s_j^{NMDA}$ and s_j^{GABA} are the receptor-specific fractions of open channels. w_j are the synaptic weights, defining the pools. The NMDA receptor-activated synaptic currents are dependent on the membrane voltage

and are controlled by the extracellular concentration of magnesium $[Mg^{2+}]$. The fraction of open channels is given by:

$$\frac{ds_j^{AMPA,ext}(t)}{dt} = -\frac{s_j^{AMPA,ext}(t)}{\tau_{AMPA}} + \sum_k \delta(t - t_j^k) \qquad (A.7)$$

$$\frac{ds_j^{AMPA,rec}(t)}{dt} = -\frac{s_j^{AMPA,rec}(t)}{\tau_{AMPA}} + \sum_k \delta(t - t_j^k) \qquad (A.8)$$

$$\frac{ds_j^{NMDA}(t)}{dt} = -\frac{s_j^{NMDA}(t)}{\tau_{NMDA,decay}} + \alpha x_j(t)(1 - s_j^{NMDA}(t)) \qquad (A.9)$$

$$\frac{dx_j^{NMDA}(t)}{dt} = -\frac{x_j^{NMDA}(t)}{\tau_{NMDA,rise}} + \sum_k \delta(t - t_j^k) \qquad (A.10)$$

$$\frac{ds_j^{GABA}(t)}{dt} = -\frac{s_j^{GABA}(t)}{\tau_{GABA}} + \sum_k \delta(t - t_j^k), \qquad (A.11)$$

where $\tau_{AMPA,ext}$, $\tau_{AMPA,rec}$, $\tau_{NMDA,decay}$ and τ_{GABA} are the decay time constants, and $\tau_{NMDA,rise}$ is the rise time constant for NMDA synapses. The rise times for AMPA and GABA synapses are neglected as they are smaller than 1 ms. The sums over k represent the sums over spikes formulated as δ-peaks ($\delta(t)$) emitted by presynaptic neuron j at time k.

The values for the neuronal and synaptic dynamics are provided in Table 4.3. We implemented synaptic dynamics, and, although not essential for the model, included in them the slow synaptic dynamics produced by the long time constant of the NMDA receptors, for as shown by Wang (1999), in order to stabilize the network the recurrent excitation should be dominated by slow synaptic dynamics, such as those produced by the NMDA receptors. Moreover we used these synaptic dynamics in order to avoid synchrony and oscillations, following Brunel and Wang (2003). We did not implement synaptic delays for simplicity. Transmission delays help to prevent synchrony in the network **?**, producing similar effects in this respect to slow synaptic dynamics (implemented in the NMDA receptors). Since we used NMDA

receptors, we expect our results to be remain valid without explicitly implementing synaptic delays.

Mean-field approximation

The essence of the mean-field approximation is to simplify the integrate-and-fire dynamics by replacing after the diffusion approximation **?**, the sums of the synaptic components by the average DC component and a fluctuation term. The stationary dynamics of each population can be described by the *population transfer function*, which provides the average population rate as a function of the average input current. The set of stationary, self-reproducing rates for the different populations in the network can be found by solving a set of coupled self-consistency equations.

The mean-field approximation assumes that the network of integrate-and-fire neurons is in a stationary state. In this formulation the potential of a neuron is calculated as:

$$\tau_x \frac{dV(t)}{dt} = -V(t) + \mu_x + \sigma_x \sqrt{\tau_x} \eta(t) \tag{B.1}$$

where $V(t)$ is the membrane potential, x labels the populations, τ_x is the effective membrane time constant, μ_x is the mean value the membrane potential would have in the absence of spiking and fluctuations, σ_x measures the magnitude of the fluctuations and η is a Gaussian process with absolute exponentially decaying correlation function with

time constant τ_{AMPA}. The quantities μ_x and σ_x^2 are given by:

$$\mu_x = \frac{(T_{ext}\nu_{ext} + T_{AMPA}n_x^{AMPA} + \rho_1 n_x^{NMDA})V_E + \rho_2 n_x^{NMDA}\langle V \rangle}{S_x} +$$

(B.2)

$$+ \frac{T_I n_x^{GABA}V_I + V_L}{S_x}$$

$$\sigma_x^2 = \frac{g_{AMPA,ext}^2(\langle V \rangle - V_E)^2 N_{ext}\nu_{ext}\tau_{AMPA}^2\tau_x}{g_m^2\tau_m^2},$$

(B.3)

where ν_{ext} Hz is the external incoming spiking rate, ν_I is the spiking rate of the inhibitory population, $\tau_m = C_m/g_m$ with the values for the excitatory or inhibitory neurons depending on the population

considered and the other quantities are given by:

$$S_x = 1 + T_{ext}\nu_{ext} + T_{AMPA}n_x^{AMPA} + (\rho_1 + \rho_2)n_x^{NMDA} + T_I n_x^{GABA}$$

$$\text{(B.4)}$$

$$\tau_x = \frac{C_m}{g_m S_x} \tag{B.5}$$

$$n_x^{AMPA} = \sum_{j=1}^{p} r_j w_{jx}^{AMPA} \nu_j \tag{B.6}$$

$$n_x^{NMDA} = \sum_{j=1}^{p} r_j w_{jx}^{NMDA} \psi(\nu_j) \tag{B.7}$$

$$n_x^{GABA} = \sum_{j=1}^{p} r_j w_{jx}^{GABA} \nu_j \tag{B.8}$$

$$\psi(\nu) = \frac{\nu\tau_{NMDA}}{1 + \nu\tau_{NMDA}} \left(1 + \frac{1}{1 + \nu\tau_{NMDA}} \sum_{n=1}^{\infty} \frac{(-\alpha\tau_{NMDA,rise})^n T_n(\nu)}{(n+1)!} \right)$$

$$\text{(B.9)}$$

$$T_n(\nu) = \sum_{k=0}^{n} (-1)^k \binom{n}{k} \frac{\tau_{NMDA,rise}(1 + \nu\tau_{NMDA})}{\tau_{NMDA,rise}(1 + \nu\tau_{NMDA}) + k\tau_{NMDA,decay}}$$

$$\text{(B.10)}$$

$$\tau_{NMDA} = \alpha\tau_{NMDA,rise}\tau_{NMDA,decay} \tag{B.11}$$

$$T_{ext} = \frac{g_{AMPA,ext}\tau_{AMPA}}{g_m} \tag{B.12}$$

$$T_{AMPA} = \frac{g_{AMPA,rec}N_E\tau_{AMPA}}{g_m} \tag{B.13}$$

$$\rho_1 = \frac{g_{NMDA}N_E}{g_m J} \tag{B.14}$$

$$\rho_2 = \beta\frac{g_{NMDA}N_E(\langle V_x \rangle - V_E)(J-1)}{g_m J^2} \tag{B.15}$$

$$J = 1 + \gamma\exp(-\beta\langle V_x \rangle) \tag{B.16}$$

$$T_I = \frac{g_{GABA}N_I\tau_{GABA}}{g_m} \tag{B.17}$$

$$\langle V_x \rangle = \mu_x - (V_{thr} - V_{reset})\nu_x\tau_x, \tag{B.18}$$

where p is the number of excitatory populations, r_x is the fraction of neurons in the excitatory x population, $\omega_{j,x}$ the weight of the connections from population x to population j, ν_x is the spiking rate of the x excitatory population, $\gamma = [Mg^{++}]/3.57$, $\beta = 0.062$ and the average membrane potential $\langle V_x \rangle$ has a value between -55 mV and -50 mV.

The spiking rate of a population as a function of the defined quantities is then given by:

$$\nu_x = \phi(\mu_x, \sigma_x), \tag{B.19}$$

where ϕ is the transduction function of population x, which gives the output rate of a population x in terms of the inputs, which in turn depend on the rates of all of the populations.

$$\phi(\mu_x, \sigma_x) = \left(\tau_{rp} + \tau_x \int_{\beta(\mu_x, \sigma_x)}^{\alpha(\mu_x, \sigma_x)} du \sqrt{\pi} \exp(u^2)[1 + \mathrm{erf}(u)] \right)^{-1} \tag{B.20}$$

$$\alpha(\mu_x, \sigma_x) = \frac{(V_{thr} - \mu_x)}{\sigma_x} \left(1 + 0.5 \frac{\tau_{AMPA}}{\tau_x} \right) + 1.03 \sqrt{\frac{\tau_{AMPA}}{\tau_x}} - 0.5 \frac{\tau_{AMPA}}{\tau_x} \tag{B.21}$$

$$\beta(\mu_x, \sigma_x) = \frac{(V_{reset} - \mu_x)}{\sigma_x} \tag{B.22}$$

with $\mathrm{erf}(u)$ the error function and τ_{rp} the refractory period which is considered to be 2 ms for excitatory neurons and 1 ms for inhibitory neurons. To solve the equations defined by (B.19) for all xs we integrate numerically (B.18) and the differential equation below, which has fixed point solutions corresponding to equations B.19:

$$\tau_x \frac{d\nu_x}{dt} = -\nu_x + \phi(\mu_x, \sigma_x). \tag{B.23}$$

For the numerical integration we used an Euler routine with a step size of 0.1.

Bibliography

Each reference indicates the pages where it appears.

Abeles, A. (1991). *Corticonics*. Cambridge University Press, New York. 44, 141

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716723. 86

Albantakis, L. and Deco, G. (2009a). The encoding of alternatives in multiple-choice decision making. *P Natl Acad Sci USA*, 25:10308–10313. 10

Albantakis, L. and Deco, G. (2009b). The encoding of alternatives in multiple-choice decision making. *Proceedings of the National Academy of Sciences*, 106(25):10308–10313. 57, 90, 95, 112, 119, 133

Albantakis, L. and Deco, G. (2012). Changes of mind in an attractor network of decision-making. *PLoS Comput Biol*, 7(6). 11, 22, 99, 109

Amit, D. J. (1989). *Modeling Brain Function. The World of Attractor Neural Networks*. Cambridge University Press, Cambridge. 16

Amit, D. J., Brunel, N., and Tsodyks, M. V. (1994). Correlations of cortical hebbian reverberations: theory versus experiment. *Journal of Neuroscience*, 14(11 Pt 1):6435–6445. 16

Angell, F. (1907). On judgments of" like" in discrimination experiments. *The American Journal of Psychology*, pages 253–260. 24

Audley, R. (1964). Decision-making. *British Medical Bulletin*, 20(1):27–31. 28, 29, 36, 135

Bair, W., Koch, C., Newsome, W., and Britten, K. (1994). Power spectrum analysis of bursting cells in area mt in the behaving monkey. *The Journal of neuroscience*, 14(5):2870–2892. 72

Baranski, J. V. and Petrusic, W. M. (1994). The calibration and resolution of confidence in perceptual judgments. *Perception & Psychophysics*, 55(4):412–428. 29

Barttfeld, P., Wicker, B., McAleer, P., Belin, P., Cojan, Y., Graziano, M., Leiguarda, R., and Sigman, M. (2013). Distinct patterns of functional brain connectivity correlate with objective performance and subjective beliefs. *Proceedings of the National Academy of Sciences*, 110(28):11577–11582. 30, 133

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., Shadlen, M. N., Latham, P. E., and Pouget, A. (2008). Probabilistic population codes for bayesian decision making. *Neuron*, 60:1142–1152. 10

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18:22–47. 58

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev*, 113:700–765. 10, 12

Bogacz, R., Usher, M., Zhang, J., and McClelland, J. L. (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philos Trans R Soc Lond B Biol Sci*, 362:16551670. 10, 14

Bollimunta, A. and Ditterich, J. (2012). Local computation of decision-relevant net sensory evidence in parietal cortex. *Cereb Cortex*, 22(4):903–917. 9, 10

Bollimunta, A., Totten, D., and Ditterich, J. (2012). Neural dynamics of choice: Single-trial analysis of decision-related activity in parietal cortex. *J Neurosci*, 32(37):12684–12701. 78

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., and Movshon, J. A. (1996a). A relationship between behavioral choice and the visual responses of neurons in macaque mt. *Vis Neurosci*, 13:87–100. 7

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., and

Movshon, J. A. (1996b). A relationship between behavioral choice and the visual responses of neurons in macaque mt. *Visual Neuroscience*, (13):1486–1510. 90

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12(12):4745–4765. 7

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1993). Responses of neurons in macaque mt to stochastic motion signals. *Vis Neurosci*, 10:1157–1169. 7

Brody, C., Hernandez, A., Zainos, A., and Romo, R. (2003a). Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cerebral Cortex*, 13:1196–1207. 47

Brody, C. D., Hernndez, A., Zainos, A., and Romo, R. (2003b). Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cerebral Cortex*, 13(11):1196–1207. 7

Brunel, N. and Wang, X.-J. (2001a). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J Comput Neurosci*, 11:63–85. 16, 21

Brunel, N. and Wang, X. J. (2001b). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of Computational Neuroscience*, 11:63–85. 59, 119, 120, 141

Brunel, N. and Wang, X.-J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. *Journal of Neurophysiology*, 90:415–430. 143

Brunton, B., Botvinick, M., and Brody, C. (2013). Rats and humans can optimally accumulate evidence for decision-making. *science*, 340(6128):95–98. 22

Churchland, A. and Ditterich, J. (2012). New advances in understanding decisions among multiple alternatives. *Current Opinion in Neurobiology*, 22:920926. 9

Churchland, A. K., Kiani, R., Chaudhuri, R., Wang, X.-J., Pouget, A., and Shadlen, M. N. (2011). Variance as a signature of neural computations during decision making. *Neuron*, 69(4):818–831. 78

Churchland, A. K., Kiani, R., and Shadlen, M. N. (2008a). Decision-making with multiple alternatives. *Nature Neuroscience*, (11):693–702. 6, 90

Churchland, A. K., Kiani, R., and Shadlen, M. N. (2008b). Decision-making with multiple alternatives. *Nat Neurosci*, 11:693–702. 9, 10

Dayan, P. and Abbott, L. F. (2001). *Theoretical Neuroscience*. MIT Press, Cambridge, MA. 141

de Lafuente, V. and Romo, R. (2006). Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proc Natl Acad Sci, USA*, 103:14266–14271. 64

Deco, G. and Marti, D. (2007). Deterministic analysis of stochastic bifurcations in multi-stable neurodynamical systems. *Biological Cybernetics*, 96:487–496. 41

Deco, G. and Rolls, E. T. (2006). A neurophysiological model of decision-making and Weber's law. *Europeanjournalof Neuroscience*, 24:901–916. 2, 40, 47, 111

Deco, G., Rolls, E. T., Albantakis, L., and Romo, R. (2013). Brain mechanisms for perceptual and reward-related decision-making. *Progress in Neurobiology*, 103:194 – 213. 17, 20

Deco, G., Rolls, E. T., and Romo, R. (2009). Stochastic dynamics as a principle of brain function. *Progress in Neurobiology*, 88:1–16. 41

Ditterich, J. (2006a). Evidence for time-variant decision making. *Eur J Neurosci*, 24:3628–3641. 13, 14

Ditterich, J. (2006b). Stochastic models of decisions about motion direction:behavior and physiology. *Neural Networks*, 19:981–1012. 13, 14

Ditterich, J. (2010). A comparison between mechanisms of multi-alternative perceptual decision making: Ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Front Neurosci*, 4:184. 11

Draper, N. R. and Smith, H. (1966). *Applied Regression Analysis*. John Wiley and Sons, New York. 80

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience*, 32(11):3612–

3628. 63, 116, 124

Drugowitsch, J. and Pouget, A. (2012). Probabilistic vs. non-probabilistic approaches to the neurobiology of perceptual decision-making. *Current Opinion in Neurobiology*, 22(6):963 – 969. 36

Felsen, G. and Mainen, Z. (2009). Motor planning in the rat superior colliculus. In *Conference Abstract: Computational and systems neuroscience*. Frontiers in Systems Neuroscience. 49

Festinger, L. (1943). Studies in decision: I. decision-time, relative frequency of judgment and subjective confidence. *Journal of Experimental Psychology*, 32:291–306. 26, 29

Foley, P. (1959). The expression of certainty. *The American Journal of Psychology*. 23

Furman, M. and Wang, X.-J. (2008). Similarity effect and optimal control of multiple-choice decision making. *Neuron*, 60:1153–1168. 57, 119

Furman, M. and Wang, X.-J. (2009). Similarity effect and optimal control of multiple-choice decision making. *Neuron*, 60:11531168. 10

Garret, H. E. (1922). A study of the relation of accuracy to speed. *Archs Psychol.*, 56:1–105. 23, 26, 29

Gold, J. and Shadlen, M. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404:390–394. 6, 7

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–574. 33, 40

Graziano, M. and Sigman, M. (2009). The spatial and temporal construction of confidence in the visual scene. *PLoS ONE*, 4(3):e4909. 2, 23, 27, 64, 133

Green, D. M., Swets, J. A., et al. (1966). *Signal detection theory and psychophysics*, volume 1. Wiley New York. 23, 30

Hampton, R. R. (2001a). Rhesus monkeys know when they can remember. *Proceedings of the National Academy of Sciences of the USA*, 98:5539–5362. 58

Hampton, R. R. (2001b). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, 98(9):5359–5362. 25

Hampton, R. R. (2001c). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, 98(9):5359–5362. 114

Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological Review*, 18(3):186. 23, 28, 36, 135

Hernández, A., Nácher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vázquez, Y., Camarillo, L., and Romo, R. (2010). Decoding a perceptual decision process across cortex. *Neuron*, 66(2):300–314. 64

Hernandez, A., Zainos, A., and Romo, R. (2002). Temporal evolution of a decision-making process in medial premotor cortex. *Neuron*, 33:959–972. 57

Hernndez, A., Zainos, A., and Romo, R. (2002). Temporal evolution of a decision-making process in medial premotor cortex. *Neuron*, 33(6):959 – 972. 8

Heuer, H. W. and Britten, K. H. (2004). Optic flow signals in extrastriate area mst: Comparison of perceptual and neuronal sensitivity. *Journal of Neurophysiology*, 91(3):1314–1326. 6

Heyes, C. (2008). Beast machines? Questions of animal consciousness. In Weiskrantz, L. and Davies, M., editors, *Frontiers of Consciousness*, chapter 9, pages 259–274. Oxford University Press, Oxford. 58

Hick, W. (1952). On the rate of gain of information. *Q J Exp Psychol*, 4:11–26. 9

Hofstadter, D. R. (2007). *I am a Strange Loop*. Basic Books, New York, NY. 57

Horwitz, G. D. and Newsome, W. T. (2001). Target selection for saccadic eye movements: Prelude activity in the superior colliculus during a direction-discrimination task. *Journal of Neurophysiology*, 86(5):2543–2558. 78

Howard, I. S., Ingram, J. N., and Wolpert, D. M. (2009). A modular planar robotic manipulandum with end-point torque control. *Journal of Neuroscience Methods*, 181(2):199–211. 115, 133

Insabato, A., Pannunzi, M., Rolls, E. T., and Deco, G. (2010). Confidence-related decision-making. *Journal of Neurophysiology*, 104:539–547. 2, 39, 64, 69, 71, 76, 114, 117

Irwin, F. W., Smith, W. A., and Mayfield, J. F. (1956). Test of two theories of decision in an "expanded judgment" situation. *Journal of Experimental Psychology*, 51:261–268. 28, 29, 35

Johnson, D. M. (1939). Confidence and speed in the two-category judgment. *Archs Psychol.*, 34:1–53. 26, 29, 30

Jones, L. M., Fontanini, A., Sadacca, B. F., Miller, P., and Katz, D. B. (2007). Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A*, 104:1877218777. 84

Jonsson, F. U., Olsson, H., and Olsson, M. J. (2005). Odor emotionality affects the confidence in odor naming. *Chemical Senses*, 30:29–35. 49

Kemere, C., Santhanam, G., Yu, B. M., Afshar, A., I, R. S., Meng, T. H., and V, S. K. (2008). Detecting neural state transitions using hidden markov models for motor cortical prostheses. *Journal Neurophysiol*, 100:24412452. 84

Kepecs, A. and Mainen, Z. F. (2012). A computational framework for the study of confidence in humans and animals. *Phil. Trans. R. Soc. B*, 367:1322–1337. 4, 23, 24, 26, 27, 29, 80, 90, 107, 114, 129, 137

Kepecs, A., Uchida, N., Zariwala, H. A., and Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455:227–231. 2, 5, 9, 25, 27, 29, 30, 31, 32, 39, 40, 43, 44, 47, 49, 51, 53, 55, 56, 57, 58, 63, 65, 68, 72, 75, 77, 80, 107, 114, 115, 130

Kiani, A. H. and Shadlen, M. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci*, 25:10420–10436. 20

Kiani, R., Hanks, T., and Shadlen, M. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci*, 28:3017–3029. 20, 133, 136

Kiani, R. and Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, 324:759–764. 2, 5, 9, 22, 27, 33, 34, 35, 36, 64, 65, 90, 91, 92, 93, 95, 96, 98, 99, 105, 112, 114, 115, 116, 118, 131, 135

Kim, J. and Shadlen, M. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature Neuroscience*,

2:176–185. 57

Koch, C. and Preuschoff, K. (2007). Betting the house on consciousness. *Nature Neuroscience*, 10(2):140–141. 1

Laming, D. R. J. (1968). Information theory of choice-reaction times. 13

Lemus, L., Hernndez, A., Luna, R., Zainos, A., Ncher, V., and Romo, R. (2007). Neural correlates of a postponed decision report. *Proceedings of the National Academy of Sciences*, 104(43):17174–17179. 9

Lo, C.-C. and Wang, X.-J. (2006). Cortico-basal ganglia circuit mechanism for a diecison threshold in reaction time tasks. *Nature Neuroscience*, (9). 99

Louie, K., Grattan, L., and Glimcher, P. (2011). Reward value-based gain control: divisive normalization in parietal cortex. *J Neurosci*, 22:10627–10639. 9

Lycan, W. G. (1997). Consciousness as internal monitoring. In Block, N., Flanagan, O., and Guzeldere, G., editors, *The Nature of Consciousness: Philosophical Debates*, pages 755–771. MIT Press, Cambridge, MA. 58

Machens, C., Romo, R., and Brody, C. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science*, 307:1121–1124. 47

Marr, D. (1982). *Vision.* Freeman, New York. 137

Marti, D., Deco, G., Del Giudice, P., and Mattia, M. (2006). Reward-biased probabilistic decision-making: mean-field predictions and spiking simulations. *Neurocomputing*, 39:1175–1178. 2, 22, 95, 112

Martí, D., Deco, G., Mattia, M., Gigante, G., and Del Giudice, P. (2008). A fluctuation-driven mechanism for slow decision processes in reverberant networks. *PLoS ONE*, 3(7):e2534. 41

Mattia, M. and Del Giudice, P. (2002). Population dynamics of interacting spiking neurons. *Physical Review E*, 66:051917. 44

Mattia, M. and Del Giudice, P. (2004). Finite-size dynamics of inhibitory and excitatory interacting spiking neurons. *Physical Review E*, 70:052903. 44

Mazurek, M. E., Roitman, J. D., Ditterich, J., and Shadlen, M. N.

(2003). A role for neural integrators in perceptual decision making. *Cerebral Cortex*, 13(11):1257–1269. 98

Metzinger, T. (2000). *Neural correlates of consciousness: Empirical and conceptual questions*. MIT press. 137

Moreno-bote, R. (2010). Decision confidence and uncertainty in diffusion models with partially correlated neuronal integrators. *Neural Computation*, (7):1786–1811. 2, 14, 15, 36, 63, 135, 136

Mountcastle, V., Steinmetz, M., and Romo, R. (1990). Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. *The Journal of Neuroscience*, 10(9):3032–3044. 7

Nienborg, H. and Cumming, B. G. (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature*, 459(7243):89–92. 6

Niwa, M. and Ditterich, J. (2008a). Perceptual decisions between multiple directions of visual motion. *J Neurosci*, 28(17):4435–4445. 10

Niwa, M. and Ditterich, J. (2008b). Perceptual decisions between multiple directions of visual motion. *The Journal of Neuroscience*, 28(17):4435–4445. 22

Pannunzi, M., Gigante, G., Mattia, M., Deco, G., Fusi, S., and Del Giudice, P. (2012). Learning selective top-down control enhances performance in a visual categorization task. *Journal of Neurophysiology*, 108(11):3124–3137. 2, 22

Pardo-Vazquez, J. L., Leboran, V., and Acua, C. (2008). Neural correlates of decisions and their outcomes in the ventral premotor cortex. *The Journal of Neuroscience*, 28(47):12396–12408. 6, 9

Pardo-Vázquez, J. L., Leboran, V., and C., A. (2008). Neural correlates of decisions and their outcomes in the ventral premotor cortex. *Journal of Neuroscience*, 28:12396–12408. 64, 65, 66, 77, 79, 80

Pardo-Vázquez, J. L., Leboran, V., and C., A. (2009). A role for the ventral premotor cortex beyond performance monitoring. *Proc Natl Acad Sci, USA*, 106:18815–18819. 64

Persaud, N. and Mcleod, P. (2008). Wagering demonstrates subconscious processing in a binary exclusion task. *Consciousness and Cognition*, 17(3):565–575. 1, 25

Pierce, C. S. and Jastrow, J. (1884). On small differences in sensetion.

Proceedings of the National Academy of Sciences of the USA, 3:73–83. 1, 23, 29, 114

Pierrel, R. and Murray, C. S. (1963). Some relation between comparative judgment confidence and decision-time in weight lifting. *American Journal of Psychology*, 76:28–38. 27, 28

Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: a theory of choice, decision time and confidence. *Psychological Review*, (3):864–901. 2, 64, 72, 76, 114

Ponce Alvarez, A., Kilavik, B., and Riehle, A. (2008). Dynamic sequences of states in ensembles of motor cortical neurons. In *Deuxième conférence française de Neurosciences Computationnelles, "Neurocomp08"*, Marseille, France. 69

Ponce-Alvarez, A., Nácher, V., Luna, ., Riehle, A., and Romo, R. (2012). Dynamics of cortical neuronal ensembles transit from decision making to storage for later report. *The Journal of Neuroscience*, 32(35):11956–11969. 78, 84, 85

Rabiner, L. (1989). A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of IEEE*, 77(2):716723. 83

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol Rev*, 85:59–108. 13, 22

Ratcliff, R., Love, J., Thompson, C., and Opfer, J. (2012). Children are not like older adults: A diffusion model analysis of developmental changes in speeded responses. *Child Development*, 83:367–381. 14

Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput*, 20:873–922. 13, 14

Ratcliff, R. and Rouder, J. F. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9:347–356. 40, 99

Ratcliff, R. and Smith, P. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev*, 111:333–367. 14

Ratcliff, R., Zandt, T. V., and McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Reviews*, 106:261–300. 40

Reed, J. B. (1951). The speed and accuracy in discriminating differences in hue, brilliance, area and shape. In Johnson, D. M., editor,

The Psychology of thought and Judgment, pages 371–372. Harper, New York. 28, 36, 135

Resulaj, A., Kiani, R., Wolpert, D. M., and Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, 461(7261):263–266. 11, 12, 108, 115, 133

Rockette, H. E., Gur, D., and Metz, C. E. (1992). The use of continuous and discrete confidence judgments in receiver operating characteristic studies of diagnostic imaging techniques. *Investigative radiology*, 27(2):169–172. 23

Roe, R. M., Busemeyer, J. R., and Townsend, J. T. (2001). Multi-alternative decision field theory: A dynamic connectionst model of decision making. *Psychological review*, 108(2):370. 9

Roitman, J. D. and Shadlen, M. N. (2002a). Responce of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, 21(22):9475–9489. 78, 90

Roitman, J. D. and Shadlen, M. N. (2002b). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci*, 22(21):94759489. 7

Rolls, E. T. (2007). A computational neuroscience approach to consciousness. *Neural Networks*, 20:962–982. 58

Rolls, E. T. and Deco, G. (2010). *The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function*. Oxford University Press, Oxford. 41, 47, 49, 57

Rolls, E. T., Grabenhorst, F., and Deco, G. (2010). Decision-making, errors, and confidence in the brain. *Submitted.* 49

Romo, R., Hernández, A., and Zainos, A. (2004). Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron*, 41:165–73. 64

Romo, R., Hernandez, A., Zainos, A., Lemus, L., and Brody, C. (2002a). Neural correlates of decision-making in secondary somatosensory cortex. *Nature Neuroscience*, 5:1217–1225. 57

Romo, R., Hernndez, A., Zainos, A., Lemus, L., and Brody, D. C. (2002b). Neuronal correlates of decision-making in secondary somatosensory cortex. *Nature Neuroscience*, 5:1217–1225. 7

Roxin, A. and Ledberg, A. (2008a). Neurobiological models of two-

choice decision making can be reduced to a one-dimensional non-linear diffusion equation. *PLoS Comput Biol*, 4(3):e1000046. 2, 21, 111, 134

Roxin, A. and Ledberg, A. (2008b). Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation. *PLoS Comput Biol*, 4. 21

Rushworth, M. F. S. and Behrens, T. E. J. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience*, 11:389–397. 2

Salinas, E., Hernndez, A., Zainos, A., and Romo, R. (2000). Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *The Journal of Neuroscience*, 20(14):5503–5515. 7

Sallet, J. and Rushworth, M. F. S. (2009). Should i stay or should i go: Genetic bases for uncertainty-driven exploration. *Nature Neuroscience*, 12:963–965. 2

Sanders, J. I. and Kepecs, A. (2012). Choice ball: a response interface for two-choice psychometric discrimination in head-fixed mice. *Journal of Neurophysiology*, 108(12):3416–3423. 115, 133

Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464. 86

Seidemann, E., Meilijson, I., Abeles, M., Bergman, H., and Vaadia, E. (1996). Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *Journal Neuroscience*, 16:752768. 84

Shadlen, M., Britten, K., Newsome, W., and Movshon, J. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *The Journal of Neuroscience*, 16(4):1486–1510. 7

Shadlen, M. and Newsome, W. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *J Neurophysiol*, 86:19161936. 6, 7, 16

Shadlen, M. N. and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current opinion in neurobiology*, 4(4):569–579. 72

Shields, W. E., Smith, J. D., Guttmannova, K., and Washburn, D. A. (2005). Confidence judgments by humans and rhesus monkeys. *Journal of General Psychology*, 132(2):165–186. 26, 27

Shields, W. E., Smith, J. D., and Washburn, D. A. (1997). Uncertain responses by humans and rhesus monkeys (macaca mulatta) in a psychophysical samedifferent task. *J. Exp. Psychol. Gen.*, 126:147–164. 24, 91, 107, 114

Simoncelli, E. P. and Heeger, D. J. (1998). A model of neuronal responses in visual area mt. *Vision Research*, (5):743–761. 90

Smith, J. D., Beran, M. J., Redford, J. S., and Washburn, D. A. (2006). Dissociating uncertainty responses and reinforcement signals in the comparative study of uncertainty monitoring. *J. Exp. Psychol. Gen.*, 135:282–297. 91, 107, 114

Smith, J. D., Schull, J., Strote, J., McGee, K., Egnor, R., and Erb, L. (1995). The uncertain response in the bottlenosed dolphin (tursiops truncatus). *Journal of Experimental Psychology: General*, 124(4):391. 24, 26

Smith, J. D., Shields, W. E., Schull, J., and Washburn, D. A. (1997). The uncertain response in humans and animals. *Cognition*, 62(1):75–97. 24

Smith, J. D., Shields, W. E., and Washburn, D. A. (2003). The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and Brain Sciences*, 26(3):317–339. 24

Smith, P. L. (2010). From poisson shot noise to the integrated ornstein–uhlenbeck process: Neurally principled models of information accumulation in decision-making and response time. *Journal of Mathematical Psychology*, 54(2):266–283. 21

Smith, P. L. and McKenzie, C. R. (2011). Diffusive information accumulation by minimal recurrent neural models of decision making. *Neural computation*, 23(8):2000–2031. 21

Snowden, R. J., Stefan, T., Erickson, R. G., and Andersen, R. A. (1991a). The response of area mt and v1 neurons to transparent motion. *Journal of Neuroscience*, (11):2768–2786. 90

Snowden, R. J., Treue, S., Erickson, R. G., and Andersen, R. A. (1991b). Response of area mt and v1 neurons to transparent motion. *J Neurosci*, 11:2768–2795. 6

Sober, E. (1994). *From a biological point of view*, chapter Let's Razor Ockham's Razor, pages 136–157. Cambridge University Press. 15, 135

Softky, W. R. and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *The Journal of Neuroscience*, 13(1):334–350. 72

Stein, R. B. (1965). A theoretical analysis of neuronal variability. *Biophysical Journal*, 5(2):173–194. 21

Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25(3):251–260. 13

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological review*, 79(4):281. 9

Tversky, A. and Simonson, I. (1993). Context-dependent preferences. *Management science*, 39(10):1179–1189. 9

Usher, M. and McClelland, J. (2001). On the time course of perceptual choice: the leaky competing accumulator model. *Psychological Reviews*, 108:550–592. 9, 40

Vandekerckhove, J. and Tuerlinckx, F. (2007). Fitting the ratcliff diffusion model to experimental data. *Psychon Bull Rev*, 14:1011–1026. 14

Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, 13(1):37–58. 10, 14

Vickers, D. (1979a). *Decision Processes in Visual Perception*. Academic Press, New York. 34, 35, 49

Vickers, D. (1979b). *Decision Processes in Visual Perception*. Academic Press, New York. 63, 69, 80, 104, 107, 116, 132, 133, 136

Vickers, D. and Packer, J. (1982). Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. *Acta Psychologica*, 50:179–197. 30, 40, 49

Vickers, D., Smith, P., Burt, J., and Brown, M. (1985). Experimental paradigms emphasising state or process limitations: Ii effects on confidence. *Acta Psychologica*, 59(2):163–193. 27

Volkmann, J. (1934). The relation of time of judgment to certainty of judgment. *Psychological Bulletin*, 31:672–673. 28, 36, 135

Vzquez, P., Cano, M., and Acua, C. (2000). Discrimination of line orientation in humans and monkeys. *Journal of Neurophysiology*, 83(5):2639–2648. 6

Wang, X. J. (1999). Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *journalof Neuroscience*, 19:9587–9603. 143

Wang, X. J. (2002a). Probabilistic decision making by slow reverberation in cortical circuit. *Neuron*, 36:955968. 16

Wang, X. J. (2002b). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36:955–968. 2, 40, 41, 47, 90, 94, 95, 112, 117

Wang, X.-J. (2012). Neural dynamics and circuit mechanisms of decision-making. *Current Opinion in Neurobiology*, (22). 99

Watson, C. S., Kellogg, S. C., Kawanishi, D. T., and Lucas, P. A. (1973). The uncertain response in detection-oriented psychophysics. *Journal of Experimental Psychology*, 99(2):180. 24

Watson, C. S., Rilling, M. E., and Bourbon, W. T. (1964). Receiver-operating characteristics determined by a mechanical analog to the rating scale. *The Journal of the Acoustical Society of America*, 36:283. 23

Weiskrantz, L. (1997). *Consciousness Lost and Found.* Oxford University Press, Oxford. 58

Wong, K. and Huk, A. (2008). Temporal dynamics underlying perceptual decision making: Insights from the interplay between an attractor model and parietal neurophysiology. *Front Neurosci*, 2:245–254. 2, 20

Wong, K., Huk, A., Shadlen, M., and Wang, X. (2007). Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front Comput Neurosci*, 1:6. 20

Wong, K. F. and Wang, X. J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci*, 26:13141328. 2, 21

Zylberberg, A., Barttfeld, P., and Sigman, M. (2012). The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience*, 6:79. 2, 133, 136