The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function

Edmund T. Rolls
Oxford Centre for Computational Neuroscience
Oxford
England

Gustavo Deco
Institucio Catalana de Recerca i Estudis Avancats (ICREA)
Universitat Pompeu Fabra
Barcelona
Spain

OXFORD UNIVERSITY PRESS • OXFORD 2010

Preface

The relatively random spiking times of individual neurons produce a source of noise in the brain. The aim of this book is to consider the effects of this and other noise on brain processing. We show that in cortical networks this noise can be an advantage, for it leads to probabilistic behaviour that is advantageous in decision-making, by preventing deadlock, and is important in signal detectability. We show how computations can be performed through stochastic dynamical effects, including the role of noise in enabling probabilistic jumping across barriers in the energy landscape describing the flow of the dynamics in attractor networks. The results obtained in neurophysiological studies of decision-making and signal detectability are modelled by the stochastical neurodynamics of integrate-and-fire networks of neurons with probabilistic neuronal spiking. We describe how these stochastic neurodynamical effects can be analyzed, and their importance in many aspects of brain function, including decision-making, perception, memory recall, short-term memory, and attention. We show how instabilities in these brain dynamics may contribute to the cognitive symptoms in aging and in psychiatric states such as schizophrenia, and how overstability may contribute to the symptoms in obsessive-compulsive disorder.

This is a new approach to the dynamics of neural processing, in which we show that noise breaks deterministic computations, and has many advantages. These principles need to be analyzed in order to understand brain function and behaviour, and it is an aim of this book to elucidate the stochastic, that is probabilistic, dynamics of brain processing, and its advantages. The book describes approaches that provide a foundation for this understanding, including integrate-and-fire models of brain and cognitive function that incorporate the stochastic spiking-related dynamics, and mean-field analyses that are consistent in terms of the parameters with these, but allow formal analysis of the networks. A feature of the treatment of the mean-field approach is that we introduce new ways in which it can be extended to include some of the effects of noise on the operation of the system. The book thus describes the underpinnings in physics of this new approach to the probabilistic functioning of the brain. However, at the same time, most of the concepts of the book and the principles of the stochastic operation of the brain described in the book, can be understood by neuroscientists and others interested in brain function who do not have expertise in mathematics or theoretical physics, and the book has been written with this in mind.

We believe that the principles of the stochastic dynamics of brain function described in this book are important, for brain function can not be understood as a deterministic noiseless system.

To understand how the brain works, including how it functions in memory, attention, and decision-making, it is necessary to combine different approaches, including neural computation. Neurophysiology at the single neuron level is needed because this is the level at which information is exchanged between the computing elements of the brain, the neurons. Evidence from the effects of brain damage, including that available from neuropsychology, is needed to help understand what different parts of the system do, and indeed what each part is necessary for. Functional neuroimaging is useful to indicate where in the human brain different processes take place, and to show which functions can be dissociated from each other. Knowledge of the biophysical and synaptic properties of neurons is essential to understand how the computing elements of the brain work, and therefore what the building blocks

of biologically realistic computational models should be. Knowledge of the anatomical and functional architecture of the cortex is needed to show what types of neuronal network actually perform the computation. The approach of neural computation is also needed, as this is required to link together all the empirical evidence to produce an understanding of how the system actually works. But an understanding of the role of noise in brain computation is also crucial, as we show in this book. This book utilizes evidence from all these approaches to develop an understanding of how different types of memory, perception, attention, and decision-making are implemented by processing in the brain, and are influenced by the effects of noise.

We emphasize that to understand memory, perception, attention, and decision-making in the brain, we are dealing with large-scale computational systems with interactions between the parts, and that this understanding requires analysis at the computational and global level of the operation of many neurons to perform together a useful function. Understanding at the molecular level is important for helping to understand how these large-scale computational processes are implemented in the brain, but will not by itself give any account of what computations are performed to implement these cognitive functions. Instead, understanding cognitive functions such as memory recall, attention, and decision-making requires single neuron data to be closely linked to computational models of how the interactions between large numbers of neurons and many networks of neurons allow these cognitive problems to be solved. The single neuron level is important in this approach, for the single neurons can be thought of as the computational units of the system, and is the level at which the information is exchanged by the spiking activity between the computational elements of the brain. The single neuron level is therefore, because it is the level at which information is communicated between the computing elements of the brain, the fundamental level of information processing, and the level at which the information can be read out (by recording the spiking activity) in order to understand what information is being represented and processed in each brain area. Moreover, the probabilistic spiking of individual neurons is an important source of noise in the brain, and must be taken into account to understand brain function.

A test of whether one's understanding is correct is to simulate the processing on a computer, and to show whether the simulation can perform the tasks of memory systems in the brain, and whether the simulation has similar properties to the real brain. The approach of neural computation leads to a precise definition of how the computation is performed, and to precise and quantitative tests of the theories produced. How memory systems in the brain work is a paradigm example of this approach, because memory-like operations which involve altered functionality as a result of synaptic modification are at the heart of how all computations in the brain are performed. It happens that attention and decision-making can be understood in terms of interactions between and fundamental operations of networks that implement computations that implement memory operations in the brain, and therefore it is natural to treat these areas of cognitive neuroscience as well as memory in this book. The same fundamental concepts based on the operation of neuronal circuitry can be applied to all these functions, as is shown in this book.

One of the distinctive properties of this book is that it links the neural computation approach not only firmly to neuronal neurophysiology, which provides much of the primary data about how the brain operates, but also to psychophysical studies (for example of attention); to psychiatric studies of patients; to functional magnetic resonance imaging (fMRI) (and other neuroimaging) approaches; and to approaches influenced by theoretical physics about how the operation of large scale systems can be understood as a result of statistical effects in its components, in this case the neurons. The empirical evidence that is brought to bear is largely from non-human primates and from humans, because of the considerable similarity of their memory and related systems, and the overall aims to understand how memory,

attention, decision-making and related functions are implemented in the human brain, and the disorders that can arise.

The overall aims of the book are developed further, and the plan of the book is described, in Chapter 1, Section 1.1.

Part of the material described in the book reflects work performed in collaboration with many colleagues, whose tremendous contributions are warmly appreciated. The contributions of many will be evident from the references cited in the text. Especial appreciation is due to Alessandro Treves, Marco Loh, and Simon M. Stringer, who have contributed greatly in an always interesting and fruitful research collaboration on computational aspects of brain function, and to many neurophysiology and functional neuroimaging colleagues who have contributed to the empirical discoveries that provide the foundation to which the computational neuroscience must always be closely linked, and whose names are cited throughout the text. Much of the work described would not have been possible without financial support from a number of sources, particularly the Medical Research Council of the UK, the Human Frontier Science Program, the Wellcome Trust, and the James S. McDonnell Foundation. The book was typeset by the Edmund Rolls using LaTex and WinEdt, and Gustavo Deco took primary responsibility for the Appendix.

The covers show part of the picture *Ulysses and the Sirens* painted in 1909 by Herbert James Draper. The version on the back cover has noise added, and might be called *Ulysses and the Noisy Sirens*. The metaphors are of noise: sirens, and stormy, irregular, water; of waves and basins of attraction: the waves on the horizon; of decision-making: the rational conscious in Ulysses resisting the gene-based emotion-related attractors; and of Ulysses the explorer (the Greek Odysseus of Homer), always and indefatigably (like the authors) seeking new discoveries about the world (and how it works).

Updates to some of the publications cited in this book are available at **http://www.oxcns.org**.

We dedicate this work to the overlapping group: our families, friends, and many colleagues whose contributions are greatly appreciated – *in salutem praesentium, in memoriam absentium*. In addition, Gustavo Deco thanks and dedicates this book to his family, Maria Eugenia, Nikolas, Sebastian, Martin, and Matthias. We remember too a close colleague and friend, the theoretical physicist Daniel Amit, who contributed much to the analysis of attractor networks (Amit 1989, Brunel and Amit 1997).

Contents

5 Probabilistic decision-making

5.1 Introduction

In this Chapter, we show how an attractor network can model probabilistic decision-making. Attractor or autoassociation memory networks that can implement short-term memory are described in Section 1.10. For decision-making, the attractor network is trained to have two (or more) attractor states, each one of which corresponds to one of the decisions. Each attractor set of neurons receives a biasing input which corresponds to the evidence in favour of that decision. When the network starts from a state of spontaneous firing, the biasing inputs encourage one of the attractors to gradually win the competition, but this process is influenced by the Poisson-like firing (spiking) of the neurons, so that which attractor wins is probabilistic. If the evidence in favour of the two decisions is equal, the network chooses each decision probabilistically on 50% of the trials. The model not only shows how probabilistic decision-making could be implemented in the brain, but also how the evidence can be accumulated over long periods of time because of the integrating action of the attractor short-term memory network; how this accounts for reaction times as a function of the magnitude of the difference between the evidence for the two decisions (difficult decisions take longer); and how Weber's law appears to be implemented in the brain. Details of the model are provided by Deco and Rolls (2006), and in Section 2.2.3 for the integrate-and-fire implementation, and Sections 2.7.2 and 2.7.3 for the mean-field implementation.

It is very interesting that the model of decision-making is essentially the same as an attractor model of long-term memory or short-term memory in which there are competing retrieval cues. This makes the approach very unifying, and elegant, and consistent with the presence of well-developed recurrent collateral excitatory connections in the neocortex which with the same architecture and functionality can be put to different uses. This provides for economy and efficiency in evolution and in the genetic prescription of a type of cortical architecture that can be used for many functions (see further Section 8.18 on page 253).

The link between perception and action can be conceptualized by a chain of neural operations, which leads a stimulus to guide behaviour to make a decision to choose a particular action or motor response. For example, when subjects discriminate two stimuli separated by a time interval, the chain of neural operations encompasses mechanisms from the encoding of sensory stimuli, the attentional filtering of relevant features, their maintenance in working memory, to the comparison leading to a motor response (Romo and Salinas 2001, Romo and Salinas 2003). The comparison step is a crucial operation in the decision-making process. A number of neurophysiological experiments on decision-making are providing information on the neural mechanisms underlying perceptual comparison, by analyzing the responses of neurons that correlate with the animal's behaviour (Werner and Mountcastle 1965, Talbot, Darian-Smith, Kornhuber and Mountcastle 1968, Salzman, Britten and Newsome 1990, Kim and Shadlen 1999, Gold and Shadlen 2000, Schall 2001, Hernandez, Zainos and Romo 2002, Romo, Hernandez, Zainos, Lemus and Brody 2002, Romo, Hernandez, Zainos and Salinas 2003, Glimcher 2003, Glimcher 2004, Romo et al. 2004, Smith and Ratcliff 2004, Sugrue, Corrado and Newsome 2005, Gold and Shadlen 2007). An impor-

tant finding is that cortical areas involved in generating motor responses also show activity reflecting a gradual accumulation of evidence for choosing one or another decision, such that the process of making a decision and action generation often can not be differentiated (see for example, Gold and Shadlen (2000) and Romo, Hernandez and Zainos (2004); but see also Section 8.12, where it is argued that different types of categorical choice take place in different cortical areas).

Complementary theoretical neuroscience models are approaching the problem by designing biologically realistic neural circuits that can perform the comparison of two signals (Wang 2002, Brody, Romo and Kepecs 2003b, Machens, Romo and Brody 2005, Wang 2008). These models involve two populations of excitatory neurons, engaged in competitive interactions mediated by inhibition, and external sensory inputs that bias this competition in favor of one of the populations, producing a binary choice that develops gradually. Consistent with the neurophysiological findings, this neurodynamical picture integrates both the accumulation of perceptual evidence for the comparison, and choice, in one unifying network. Even more, the models are able to account for the experimentally measured psychometric and neurometric curves, and reaction times (Marti et al. 2008).

The comparison of two stimuli for which is the more intense becomes more difficult as they become more similar. The 'difference-threshold' or 'just-noticeable difference' (jnd) is the amount of change needed for us to recognize that a change has occurred. Weber's law (enunciated by Ernst Heinrich Weber 1795–1878) states that the ratio of the difference-threshold to the background intensity is a constant. Most theoretical models of decision-making (Wang 2002, Brody et al. 2003b, Machens et al. 2005) have not investigated Weber's law, and therefore a thorough understanding of the neural substrate underlying the comparison process has been missing until recently (Deco and Rolls 2006).

In this Chapter, the neurodynamical mechanisms engaged in the process of comparison in a decision-making paradigm is investigated, and these processes are related to the psychophysics, as described for example by Weber's law. That is, the probabilistic behaviour of the neural responses responsible for detecting a just noticeable stimulus difference is part of what is described.

5.2 Decision-making in an attractor network

Let us consider the attractor network architecture again, but this time as shown in Fig. 5.1a with two competing inputs λ_1 and λ_2, each encouraging the network to move from a state of spontaneous activity into the attractor corresponding to λ_1 or to λ_2. These are separate attractor states that have been set up by associative synaptic modification, one attractor for the neurons that are coactive when λ_1 is applied, and a second attractor for the neurons that are coactive when λ_2 is applied. When λ_1 and λ_2 are both applied simultaneously, each attractor competes through the inhibitory interneurons (not shown), until one wins the competition, and the network falls into one of the high firing rate attractors that represents the decision. The noise in the network caused by the random spiking of the neurons means that on some trials, for given inputs, the neurons in the decision 1 attractor are more likely to win, and on other trials the neurons in the decision 2 attractor are more likely to win. This makes the decision-making probabilistic, for, as shown in Fig. 5.1b, the noise influences when the system will jump out of the spontaneous firing stable (low energy) state S, and whether it jumps into the high firing state for decision 1 or decision 2 (D).

The operation and properties of this model of decision-making (Wang 2002, Deco and Rolls 2006) are described first in this Chapter (5), and then in Section 5.8 we build on what

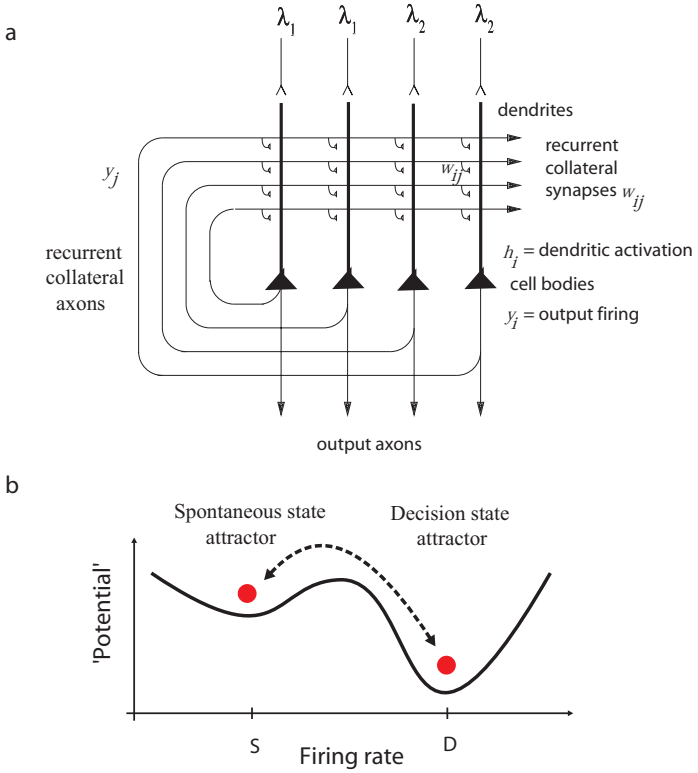

Fig. 5.1 (a) Attractor or autoassociation network architecture for decision-making. The evidence for decision 1 is applied via the λ_1 inputs, and for decision 2 via the λ_2 inputs. The synaptic weights w_{ij} have been associatively modified during training in the presence of λ_1 and at a different time of λ_2. When λ_1 and λ_2 are applied, each attractor competes through the inhibitory interneurons (not shown), until one wins the competition, and the network falls into one of the high firing rate attractors that represents the decision. The noise in the network caused by the random spiking of the neurons means that on some trials, for given inputs, the neurons in the decision 1 attractor are more likely to win, and on other trials the neurons in the decision 2 attractor are more likely to win. This makes the decision-making probabilistic, for, as shown in (b), the noise influences when the system will jump out of the spontaneous firing stable (low energy) state S, and whether it jumps into the high firing state for decision 1 or decision 2 (D).

has been described, and introduce the different stable states of this type of model, and the effects of noise in such a system.

5.3 The neuronal data underlying a vibrotactile discrimination task

A good paradigm for studying the mechanisms of decision-making is the vibrotactile sequential discrimination task, because evidence on the neuronal basis is available. In the two-alternative, forced-choice task used, subjects must decide which of two mechanical vibrations applied sequentially to their fingertips has the higher frequency of vibration (see Fig. 7.1 on page 214). Neuronal recording and behavioural analyses (Romo and Salinas 2003) have provided sufficient detail about the neuronal bases of these decisions for a quantitative model to be developed. In particular, single neuron recordings in the ventral premotor cortex (VPC) reveal neurons whose firing rate was dependent only on the difference between the two applied

frequencies, the sign of that difference being the determining factor for correct task performance (Romo, Hernandez and Zainos 2004, de Lafuente and Romo 2005). We consider other potential applications of this approach to decision-making in Chapter 8.

Deco and Rolls (2006) modelled the activity of these VPC neurons by means of a theoretical framework first proposed by Wang (2002), but investigated the role of finite-size fluctuations in the probabilistic behaviour of the decision-making neurodynamics, and especially the neural encoding of Weber's law. (The finite size fluctuations are the statistical effects caused by the chance firing of different numbers of spiking neurons in a short time period, which are significant in networks of less that infinite size, and which can influence the way in which a network operates or settles, as described further below.) The neurodynamical formulation is based on the principle of biased competition/cooperation that has been able to simulate and explain, in a unifying framework, attention, working memory, and reward processing in a variety of tasks and at different cognitive neuroscience experimental measurement levels (Rolls and Deco 2002, Deco and Lee 2002, Corchs and Deco 2002, Deco, Pollatos and Zihl 2002, Corchs and Deco 2004, Deco and Rolls 2002, Deco and Rolls 2003, Deco and Rolls 2004, Deco, Rolls and Horwitz 2004, Szabo, Almeida, Deco and Stetter 2004, Deco and Rolls 2005b, Deco and Rolls 2005a, Rolls and Deco 2006, Rolls, Loh, Deco and Winterer 2008b, Rolls, Grabenhorst and Deco 2010b).

The neuronal substrate of the ability to discriminate two sequential vibrotactile stimuli has been investigated by Romo and colleagues (Romo and Salinas 2001, Hernandez, Zainos and Romo 2002, Romo, Hernandez, Zainos, Lemus and Brody 2002, Romo, Hernandez, Zainos and Salinas 2003, Romo and Salinas 2003, Romo, Hernandez and Zainos 2004, de Lafuente and Romo 2005). They used a task where trained macaques (*Macaca mulatta*) must decide and report which of two mechanical vibrations applied sequentially to their fingertips has the higher frequency of vibration by pressing one of two pushbuttons. This decision-making paradigm requires therefore the following processes: (1) the perception of the first stimulus, a 500 ms long vibration at frequency f1; (2) the storing of a trace of the f1 stimulus in short-term memory during a delay of typically 3 s; (3) the perception of the second stimulus, a 500 ms long vibration at frequency f2; and (4) the comparison of the second stimulus f2 to the trace of f1, and choosing a motor act based on this comparison (f2-f1). The vibrotactile stimuli f1 and f2 utilized were in the range of frequencies called *flutter*, i.e. within approximately 5–50 Hz.

Deco and Rolls (2006) were particularly interested in modelling the responses of ventral premotor cortex (VPC) neurons (Romo, Hernandez and Zainos 2004). The activity of VPC neurons reflects the current and the remembered sensory stimuli, their comparison, and the motor response, i.e the entire cascade of decision-making processing linking the sensory evaluation to the motor response. Many VPC neurons encode f1 during both the stimulus presentation and the delay period. During the comparison period, the averaged firing rate of VPC neurons after a latency of a few hundred milliseconds reflects the result of the comparison, i.e. the sign of (f2-f1), and correlates with the behavioural response of the monkey. In particular, we are interested in VPC neurons which show the strongest response only during the comparison period and reflect the sign of the comparison f2-f1, i.e. these neurons are only activated during the presentation of f2, with some responding to the condition f1$<$f2, and others to the condition f1$>$f2. These neurons, which are shown in figure 2(G, H, I) of Romo, Hernandez and Zainos (2004) (see example in Fig. 5.2), reflect the decision-making step of the comparison, and therefore we will describe here their probabilistic dynamical behaviour as reported by the experimental work; and through theoretical analyses we will relate their behaviour to how decisions are taken in the brain, and to Weber's law.

Earlier brain areas provide inputs useful to the VPC. In the primary somatosensory area S1 the average firing rates of neurons in S1 convey information about the vibrotactile freq-

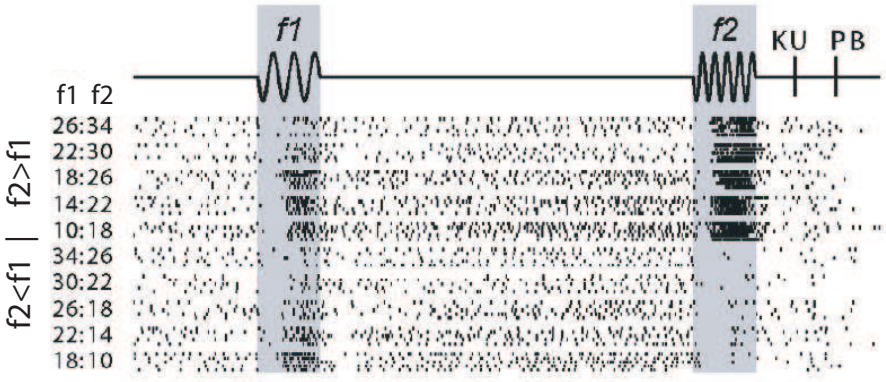

Fig. 5.2 A neuron in the ventral premotor cortex that responded to a vibrotactile stimulus during the f1 stimulation and during the delay period. However, the strongest response was for condition f2 > f1 during the f2 period. In the raster plot, each row of ticks is a trial, and each tick is an action potential. Trials were delivered in random order with 5 trials per stimulus pair shown. The labels on the left show the f1, f2 stimulus pairs for each set of 5 trials illustrated. At KU (key up) the monkey releases the key and presses either a lateral or a medial push button (PB) to indicate whether the comparison frequency (f2) was higher or lower than the base frequency (f1). (After Romo, Hernandez and Zainos 2004.)

uency f1 or f2 during the stimulation period. (The neuronal responses stop reflecting information about the stimuli immediately after the end of the stimulus.) The firing rates increase monotonically with stimulus frequency (Romo and Salinas 2003). Neurons in the secondary somatosensory area S2 respond to f1 and show significant delay activity for a few hundred milliseconds after the end of the f1 stimulus (Romo et al. 2002). Some neurons have positive and others negative monotonic relationships between their firing rate and the vibrotactile stimulus frequency. During the initial part of f2 (ca. 200 ms) the firing rate reflects either f1 or f2; later, during the last 300 ms, the firing rate reflects the comparison (f2-f1), and therefore the result of the decision. Prefrontal cortex (PFC) neurons (Brody et al. 2003a) also have a positive or negative monotonic firing rate relationship with f1. Furthermore, PFC neurons convey information about f1 into the delay period, with some neurons carrying it only during the early part of the delay period (*early neurons*), others only during the late part of the delay period (*late neurons*), and others persistently throughout the entire delay period (*persistent neurons*). During the presentation of the second stimulus f2, PFC neurons also respond like S2 neurons. Some PFC neurons respond as a function of f2 during the initial part of the comparison period, whereas other neurons show a firing rate dependency only on f1 before and at the onset of the second stimulus. In the latter part of the comparison, the firing rate reflects the comparison f2-f1. Medial premotor cortex (MPC) neurons respond similarly to PFC neurons, i.e. MPC neurons respond during f1 itself, with either positive or negative monotonic firing rate relationships, during the late part of the delay period in an f1-dependent manner in the same way as the *late* PFC neurons, and during the comparison period reflecting the comparison f2-f1 (Hernandez et al. 2002).

In summary, in the sequential vibrotactile discrimination task, S1 is predominantly sensory and the primary motor cortex (M1) is predominantly motor. A number of other cortical areas have activity that reflects the encoding, short-term memory, and comparison functions involved, perhaps as a result of information exchange between these cortical areas: the differences between S2, MPC and VPC are reflected mainly in their different latencies. In a detection task, the activity of S1 neurons codes for the stimulus but not for the behavioural choice made, whereas neuronal activity in MPC correlates with behavioural choice and de-

tection (de Lafuente and Romo 2005). Within this context, VPC (and MPC) neurons seem to reflect the core of the processing that links sensory information with action, and therefore they may represent the decision-making process itself, rather than the representation of the stimulus. Consequently VPC neurons are excellent candidates for encoding also the probabilistic behavioural response as expressed in Weber's law.

Key questions are how ventral premotor cortex (VPC) neurons (or neurons with similar activity in connected areas such as MPC) implement the decision-making process. What are the principles by which the probabilistic decisions are taken? How is the processing implemented by the neurons?

5.4 Theoretical framework: a probabilistic attractor network

The theoretical framework within which the new model was developed was utilized by Wang (2002), which is based on a neurodynamical model first introduced by Brunel and Wang (2001), and which has been recently extended and successfully applied to explain several experimental paradigms (Rolls and Deco 2002, Deco and Rolls 2002, Deco and Rolls 2003, Deco and Rolls 2004, Deco, Rolls and Horwitz 2004, Szabo, Almeida, Deco and Stetter 2004, Deco and Rolls 2005b, Rolls, Grabenhorst and Deco 2010b, Rolls, Grabenhorst and Deco 2010c). In this framework, we model probabilistic decision-making by a single attractor network of interacting neurons organized into a discrete set of populations, as depicted in Fig. 5.3. Populations or pools of neurons are defined as groups of excitatory or inhibitory neurons sharing the same inputs and connectivities. The network contains N_{E} (excitatory) pyramidal cells and N_{I} inhibitory interneurons. In the simulations, we used $N_{\mathrm{E}} = 800$ and $N_{\mathrm{I}} = 200$, consistent with the neurophysiologically observed proportion of 80% pyramidal cells versus 20% interneurons (Abeles 1991, Rolls and Deco 2002). The neurons are fully connected (with synaptic strengths as specified later).

The specific populations have specific functions in the task. In our minimal model, we assumed that the specific populations encode the categorical result of the comparison between the two sequentially applied vibrotactile stimulation, f1 and f2, i.e. the result that f1>f2 or the result that f1<f2. Each specific population of excitatory cells contains rN_{E} neurons (in our simulations $r = 0.1$). In addition there is one non-specific population, named 'Non-specific', which groups all other excitatory neurons in the modelled brain area not involved in the present task, and one inhibitory population, named 'Inhibitory', grouping the local inhibitory neurons in the modelled brain area. The latter population regulates the overall activity and implements competition in the network by spreading a global inhibition signal.

Because we were mainly interested in the non-stationary probabilistic behaviour of the network, the proper level of description at the microscopic level is captured by the spiking and synaptic dynamics of one-compartment *Integrate-and-Fire* (IF) neuron models (see Section 2.2). At this level of detail the model allows the use of realistic biophysical time constants, latencies and conductances to model the synaptic current, which in turn allows a thorough study of the realistic time scales and firing rates involved in the time evolution of the neural activity. Consequently, the simulated neuronal dynamics, that putatively underly cognitive processes, can be quantitatively compared with experimental data. For this reason, it is very useful to include a thorough description of the different time constants of the synaptic activity. The IF neurons are modelled as having three types of receptor mediating the synaptic currents flowing into them: AMPA, NMDA (both activated by glutamate), and GABA receptors. The excitatory recurrent postsynaptic currents (EPSCs) are considered to be mediated by AMPA (fast) and NMDA (slow) receptors; external EPSCs imposed onto the network

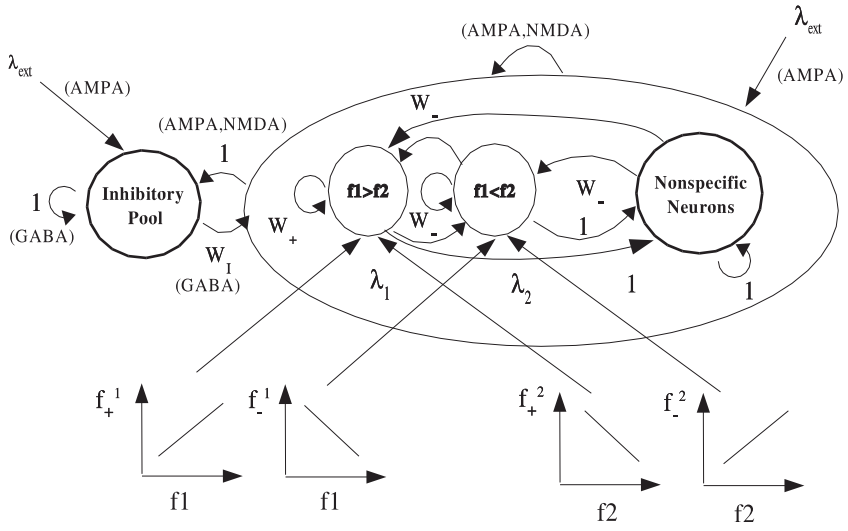

Fig. 5.3 The architecture of the neurodynamical model for a probabilistic decision-making network. The single attractor network has two populations or pools of neurons (f1>f2) and (f1<f2) which represent the decision states. One of these pools becomes active when a decision is made. If pool (f1>f2) is active, this corresponds to the decision that stimulus f1 is greater than stimulus f2. There is also a population of non-specific excitatory neurons, and a population of inhibitory neurons. Pool (f1>f2) is biased by λ_1 which reflects the strength of stimulus f1, and pool (f2>f1) is biased by λ_2 which reflects the strength of stimulus f2. (In the simulations performed f1 is the frequency of vibrotactile stimulus 1, f2 is the frequency of vibrotactile stimulus 2, and the stimuli must be compared to decide which is the higher frequency.) The integrate and fire network is subject to finite size noise, and therefore probabilistically settles into either an attractor with the population (f1>f2) active, or with the population (f1<f2) active, depending on the biasing inputs λ_1 and λ_2. The network is thus a biased competition model of decision-making. The weights connecting the different populations of neurons are shown as w_+, w_-, $w_{\rm I}$, and 1, and the values found in the mean field analysis are given in the text. All neurons receive a small random Poisson set of input spikes $\lambda_{\rm ext}$ from other neurons in the system. The nonspecific excitatory neurons are connected to pool (f1>f2) as well as to pool (f1<f2). (After Deco and Rolls 2006.)

from outside are modelled as being driven only by AMPA receptors. Inhibitory postsynaptic currents (IPSCs) to both excitatory and inhibitory neurons are mediated by GABA receptors. The details of the mathematical formulation are summarized in previous publications (Brunel and Wang 2001, Deco and Rolls 2005b, Deco and Rolls 2006), and are provided in Section 2.2.3.

We modified the conductance values for the synapses between pairs of neurons by synaptic connection weights, which can deviate from their default value 1. The structure and function of the network was achieved by differentially setting the weights within and between populations of neurons. The labelling of the weights is defined in Fig. 5.3. We assumed that the connections are already formed, by for example earlier self-organization mechanisms, as if they were established by Hebbian learning, i.e. the coupling will be strong if the pair of neurons have correlated activity (i.e. covarying firing rates), and weak if they are activated in an uncorrelated way. We assumed that the two decisions 'f1>f2' and 'f1<f2', corresponding to the two categories, are already encoded, in the sense that the monkey is already trained that pushing one or the other button, but not both, might produce a reward. As a consequence of this, neurons within a specific excitatory population are mutually coupled with a strong weight w_+, and each such population thus forms an attractor. Furthermore, the populations encoding these two decisions are likely to have anti-correlated activity in this behavioural

context, resulting in weaker than average connections between the two different populations. Consequently, we choose a weaker value $w_- = 1 - r(w_+ - 1)/(1 - r)$, so that the overall recurrent excitatory synaptic drive in the spontaneous state remains constant as w_+ is varied (Brunel and Wang 2001). Neurons in the inhibitory population are mutually connected with an intermediate weight $w = 1$. They are also connected with all excitatory neurons in the same layer with the same intermediate weight, which for excitatory-to-inhibitory connections is $w = 1$, and for inhibitory-to-excitatory connections is denoted by a weight w_{I}. Neurons in a specific excitatory population are connected to neurons in the non-selective population in the same layer with a feedforward synaptic weight $w = 1$ and a feedback synaptic connection of weight w_-.

Each individual population is driven by two different kinds of input. First, all neurons in the model network receive spontaneous background activity from outside the module through N_{ext}=800 external excitatory connections. Each connection carries a Poisson spike train at a spontaneous rate of 3 Hz, which is a typical spontaneous firing rate value observed in the cerebral cortex. This results in a background external input with a rate summed across all 800 external synapses onto each neuron of 2.4 kHz for each neuron. Second, the neurons in the two specific populations additionally receive external inputs encoding stimulus-specific information. They are assumed to originate from the somatosensory area S2 and from the PFC, encoding the frequency of both stimuli f1 (stored) and f2 (present) to be compared during the comparison period, i.e. when the second stimulus is applied. (Stimuli f1 and f2 influence λ_1 and λ_2 as shown in Fig. 5.3. The way in which the different S2 and PFC neurons described by Romo et al. (2004) are combined linearly to produce λ_1 and λ_2 is described by Deco and Rolls (2006).) These inputs which convey the evidence for each of the decisions are added to the background external inputs being applied via the 800 synapses to each neuron.

In summary, f1 is the frequency of vibrotactile stimulus 1, f2 is the frequency of vibrotactile stimulus 2, and the stimuli must be compared to decide which is the higher frequency. The single attractor network has two populations or pools of neurons (f1>f2) and (f1<f2) which represent the decision states (see Fig. 5.3). One of these pools becomes active when a decision is made. If pool (f1>f2) is active, this corresponds to the decision that stimulus f1 is greater than stimulus f2. Pool (f1>f2) is biased by λ_1 which reflects the frequency of stimulus f1, and pool (f2>f1) is biased by λ_2 which reflects the frequency of stimulus f2. The integrate and fire network is subject to finite size noise, and therefore probabilistically settles into an attractor either with the population (f1>f2) active, or with the population (f1<f2) active, depending on the biasing inputs λ_1 and λ_2. The network is thus a biased competition model of decision-making. All neurons receive a small random Poisson set of input spikes λ_{ext} via their 800 synapses that receive external inputs.

5.5 Stationary multistability analysis: mean-field

A first requirement for using the network described earlier as a probabilistic decision-making neurodynamical framework is to tune its connectivity such that the network operates in a regime of multistability. This means that at least for the stationary conditions, i.e. for periods after the dynamical transients, different possible attractors are stable. The attractors of interest for our task correspond to high activity (high spiking rates) or low activity (low spiking rates) of the neurons in the specific populations (f1>f2) and (f1<f2). The activation of the specific population (f1>f2) and the simultaneous lack of activation of the specific population (f1<f2) corresponds to encoding associated with a motor response of the monkey reporting the categorical decision f1>f2. The opposite decision corresponds to the opposite attractor states in the two specific neuronal populations. Low activity in both specific populations (the

'spontaneous state') corresponds to encoding that no decision has been made, i.e. the monkey does not answer or generates a random motor response. The same happens if both specific populations are activated (the 'pair state'). Because the monkey responds in a probabilistic way depending on the different stimuli, the operating working point of the network should be such that both possible categorical decisions, i.e. both possible single states, and sometimes (depending on the stimuli) the pair and spontaneous states, are possible stable states.

The network's operating regimes just described can all occur if the synaptic connection weights are appropriate. To determine the correct weights a mean field analysis was used (Deco and Rolls 2006) as described in Section 2.7.2 on page 85. Although a network of integrate-and-fire neurons with randomness in the spikes being received is necessary to understand the dynamics of the network, and how these are related to probabilistic decision-making, this means that the spiking activities fluctuate from time-point to time-point and from trial to trial. Consequently, integrate-and-fire simulations are computationally expensive and their results probabilistic, which makes them rather unsuitable for systematic parameter explorations. To solve this problem, we simplified the dynamics via the *mean-field* approach at least for the stationary conditions, i.e. for periods after the dynamical transients, and then analyzed the bifurcation diagrams of the dynamics. The essence of the mean-field approximation is to simplify the integrate-and-fire equations by replacing after the diffusion approximation (Tuckwell 1988, Amit and Brunel 1997, Brunel and Wang 2001), the sums of the synaptic components by the average D.C. 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 ν_i for the different populations i in the network can be found by solving a set of coupled self-consistency equations using the formulation derived by Brunel and Wang (2001) (see Section 2.7.2). The equations governing the activities in the mean-field approximation can hence be studied by standard methods of dynamical systems. The formulation departs from the equations describing the dynamics of one neuron to reach a stochastic analysis of the mean-first passage time of the membrane potentials, which results in a description of the population spiking rates as functions of the model parameters, in the limit of very large N. Obtaining a mean-field analysis of the stationary states that is consistent with the network when operating dynamically as an integrate-and-fire network is an important part of the approach used by Deco and Rolls (see Sections 2.6 and 2.7).

To investigate how the stable states depend on the connection parameters w_+ and w_{I}, Deco and Rolls (2006) solved the mean-field equations for particular values of these parameters starting at different initial conditions. For example, to investigate the stability of the state described by population (f1>f2) being in an active state and all other populations inactive, we initialize the system with that population at 10 Hz, all other excitatory populations (including the non-specific ones) at 3 Hz, and the inhibitory population at 9 Hz. If and only if, after solving the equations numerically[1], the population (f1>f2) is still active (meaning that they have a firing rate $\geq$10 Hz) but no other excitatory population is active, we conclude that the state is stable. This procedure is then repeated for all other combinations of w_+ and w_{I} to find the region where the active population (f1>f2) is stable. The stable regions of the other states are found in the same way.

Figure 5.4 presents the bifurcation diagrams resulting from the mean-field analysis, for a particular case where the behavioural decision-making is hardest and is in fact purely random (i.e. at chance) as f1 and f2 are equal. Figure 5.4 shows how the stable states of average firing vary as a function of the strength of w_+ and w_{I} for the case: f1=f2=17.5 Hz corresponding to a

[1]For all simulation periods studied, the mean-field equations were integrated using the Euler method with step size 0.1 and 4000 iterations, which always allowed for convergence.

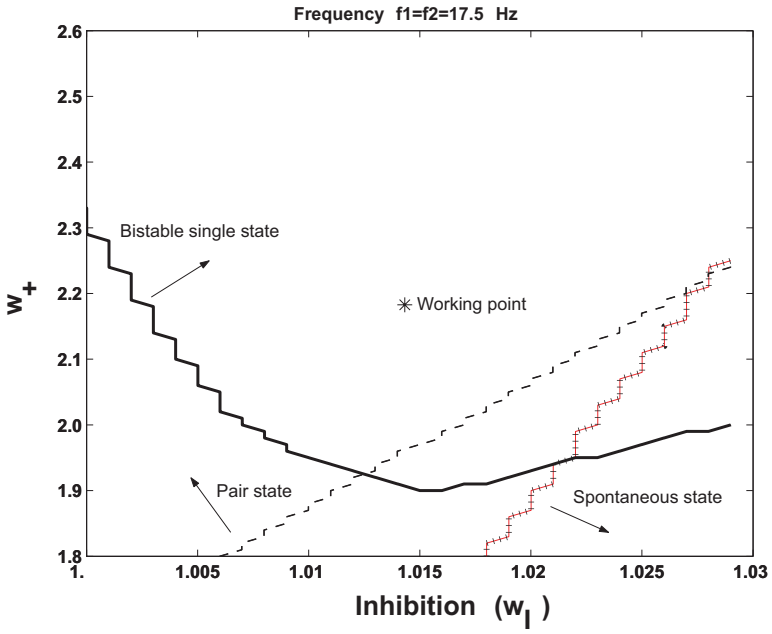

Fig. 5.4 Mean-field analysis to determine suitable values of the synaptic weights for the decision-making network. The bifurcation diagram is for the particular case where the behavioural decision-making is at chance due to f1 and f2 being equal. The diagram shows how the stable states of the average firing rate vary as a function of the synaptic strengths w_+ and w_{I} for the case: f1=f2=17.5 Hz corresponding to a low frequency of vibrotactile stimulation. The different regions where single states, a pair state, and a spontaneous firing rate state are stable are shown. In the following simulations we focus on the region of multistability (i.e. where either one or the other pool of neurons wins the competition, but where the spontaneous firing state is also a stable state), so that a probabilistic decision is possible, and therefore a convenient working point is one corresponding to a connectivity given by w_+=2.2 and w_{I}=1.015. (After Deco and Rolls 2006.)

low frequency of vibrotactile stimulation. In these cases, the specific populations (f1>f2) and (f1<f2) received an extra stimulation of λ_1 and λ_2, respectively, encoding the two vibrotactile stimuli to be compared. (The way in which these λ values were calculated simply reflects the neurons recorded in the VPC and connected areas, as described by Deco and Rolls (2006).) The different regions where single states, a pair state, and a spontaneous state are stable are shown. In the simulations, Deco and Rolls (2006) focused on a region of multistability, in which both the possible decision states, and the spontaneous firing state, were stable (see further Section 5.8), so that a probabilistic decision is possible, and therefore a convenient working point (see Fig. 5.4) is one corresponding to a connectivity given by $w_+ = 2.2$ and $w_{\mathrm{I}} = 1.015$.

Overall, Fig. 5.4 shows very large regions of stability, so that the network behaviour described here is very robust.

5.6 Non-stationary probabilistic analysis: spiking dynamics

5.6.1 Integrate-and-fire simulations of decision-making

A full characterization of the dynamics, and especially of its probabilistic behaviour, including the non-stationary regime of the system, can only be obtained through computer simulations of the spiking network model. Moreover, these simulations enable comparisons between the model in which spikes occur and neurophysiological data. The simulations of the spiking dynamics of the network were integrated numerically using the second order Runge–Kutta method with a step size 0.05 ms. Each simulation was started by a period of 500 ms where no stimulus was presented, to allow the network to stabilize. The non-stationary evolution of spiking activity was averaged over 200 trials initialized with different random seeds. In all cases, Deco and Rolls (2006) aimed to model the behaviour of the VPC neurons as shown in Fig. 5.2 from Romo, Hernandez and Zainos (2004) which reflect the decision-making performed during the comparison period. Therefore, Deco and Rolls (2006) studied the non-stationary probabilistic behaviour of the spiking network defined in Fig. 5.3, during this comparison period (during the presentation of f2), by stimulating the network simultaneously with f1 and f2. This was done by increasing the rate of the Poisson train to the neurons of both specific populations (f1>f2) and (f1<f2) by an extra value of λ_1 and λ_2, respectively, as these encode the two vibrotactile stimuli to be compared.

5.6.2 Decision-making on a single trial

Figure 5.5 shows for a single trial a typical time course of the network of VPC neurons during the decision period when the two stimuli are being compared for the case of f1=35 Hz and f2=25 Hz. The top part of Fig. 5.5 plots the time course of the mean firing rate of the populations (f1>f2), (f1<f2), and the inhibitory population. The bin widths used for the simulations were 20 ms. The transition shown corresponds to a correct trial, i.e. a transition to the correct final attractor encoding the result of the discrimination f1>f2. We observe that after 200 ms the populations (f1>f2) and (f1<f2) start to separate in such a way that the population (f1>f2) wins the competition and the network performs a transition to a single-state final attractor corresponding to a correct discrimination (i.e. high activity in the population (f1>f2) and low activity in the population (f1<f2)). The bottom part of Fig. 5.5 plots the corresponding rastergrams of 10 randomly selected neurons for each pool in the network. Each vertical line corresponds to the generation of a spike. The spiking activity shows how the firing makes the transition to the correct final single-state attractor. Further examples of the neuronal decision-making process on individual trials are shown in Figs. 5.6 and 6.2.

Romo et al. (2004) analysed the responses of VPC neurons performing the comparison as a function of both f1 and f2 (and showed the neurophysiological findings in their Figure 2G,H,I, see Fig. 5.2). Figure 5.6 shows the simulation results that correspond to those cases. Figure 5.6A shows rastergrams of a single neuron of the population (f1<f2) during the decision period starting when the f1 and f2 stimuli are applied at time = 0 ms. (The neuron should thus fire at a high rate only on trials when f2 is greater than f1, and data are shown only for correct trials.) The labels on the left indicate the vibrotactile frequencies of the f1, f2 stimulus pairs between which a decision was being made. Each row of ticks is a trial (with 10 trials of the single neuron shown for each case), and each tick is a spike. All neurons were tested with 10 trials per stimulus pair, selecting only trials where the network performed correctly. The top 5 cases correspond to a situation where f1<f2 and therefore the population (f1<f2) is highly activated after 100–200 ms. The lower 5 cases correspond to a situation where f1>f2 and therefore the population (f1<f2) is not activated. (The population (f1>f2), not shown in

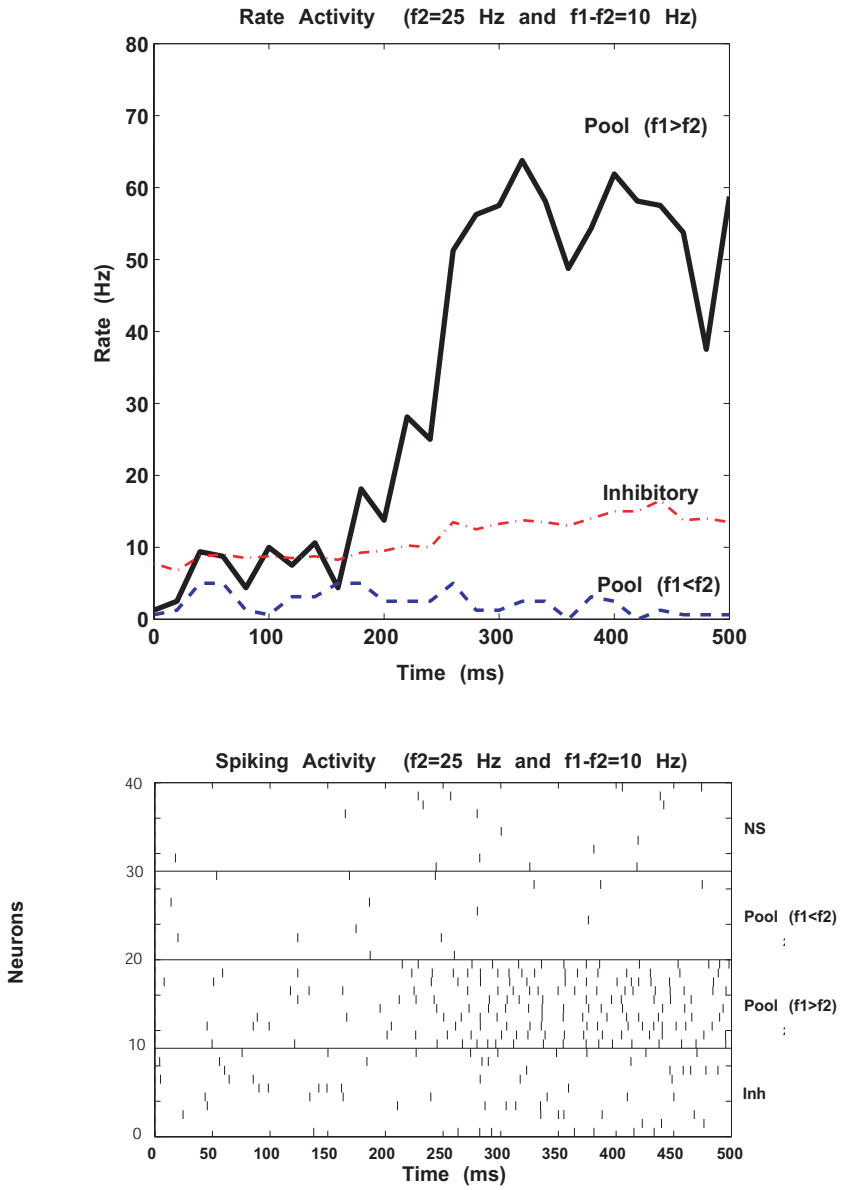

Fig. 5.5 The time course of a decision. The activity of the simulated ventral premotor cortex (VPC) neurons is shown during the decision period for the case of f1=35 Hz and f2=25 Hz. The stimuli were applied continuously starting at time = 0 ms. The top part plots the time course of the average spiking rate of the populations (f1>f2), (f1<f2), and the inhibitory population. The bin widths used for the simulations were 20 ms. The bottom part plots the corresponding rastergrams of 10 randomly selected neurons for each pool in the network. Each vertical line corresponds to the generation of a spike. In the rastergram, NS refers to neurons in the non-specific population, and Inh to neurons in the inhibitory population. The spiking activity shows the transition to the correct final single-state attractor, i.e. a transition to the correct final attractor encoding the result of the discrimination (f1>f2) (see text). (After Deco and Rolls 2006.)

Fig. 5.6, won the competition in the lower 5 cases shown in Fig. 5.6A, and therefore inhibited the (f1<f2) population in the correctly discriminated trials selected for this Figure.) Figure

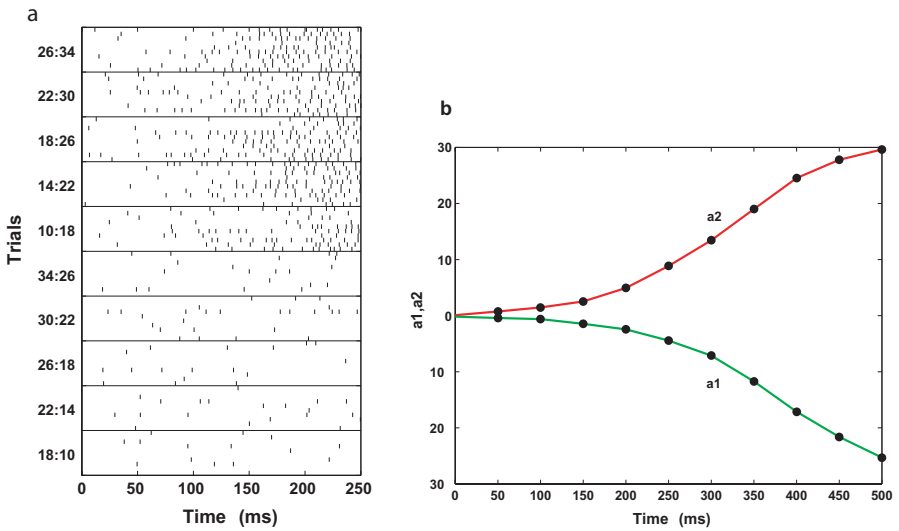

Fig. 5.6 Responses of a single neuron of the population (f1<f2) during the decision period between f1 and f2. The simulations corresponds to the experimental cases measured and studied by Romo et al. (2004) (see figure 2G,H,I of that paper, and Fig. 5.2). The different stimulation cases labelled on the left indicate f1, f2 pairs of the vibrotactile stimulation frequencies. (a) Rastergrams. Each row of ticks is a trial, and each tick is a spike. Ten correct trials are shown for each stimulus pair case. (b) Evolution of fitting coefficients $a1(t)$ and $a2(t)$ of the firing rates (see text) during the decision period. Both coefficients evolve in an antisymmetrical fashion, indicating that the average firing rate $r(t)$ is dependent primarily on the sign of the difference between f1 and f2 (i.e. $a1(t) = -a2(t)$), and not on the magnitude of the difference. (After Deco and Rolls 2006.)

5.6A shows that the firing rate at the end of the simulation period shown, at 250 ms, is high only when f2 is greater than f1, and that the final firing rate on these correct trials is relatively independent of the exact values of f1 and f2 (including the difference between f2 and f1), and responds just depending on the sign of the difference between f2 and f1, reaching a high rate only when f2 is greater than f1.

To analyse more quantitatively the dependence of the average firing rate of the VPC neurons encoding the comparisons as a function of f1 and f2, Romo et al. (2004) fitted and plotted the time evolution of the coefficients that directly measure the dependence on f1 and f2. Let us denote by $r(t)$ the trial averaged firing rate of the population (f1<f2) at time t. We determined the coefficients $a1(t)$, $a2(t)$ and $a3(t)$ that fit the firing rate of the population (f1<f2) according to $r(t) = a1(t)\text{f1} + a2(t)\text{f2} + a3(t)$. We considered for all 10 f1,f2 pair cases shown in Fig. 5.6A, 50 correct trials, using a bin width of 50 ms. Figure 5.6B plots the timecourse of $a1(t)$ and $a2(t)$ during the comparison period. In this Figure, both coefficients evolve in an antisymmetrical fashion, indicating that the average firing rate $r(t)$ depends primarily on the difference between f1 and f2 (i.e. $a1(t) = -a2(t)$) (see further Chapter 6). Even more, in the range of vibrotactile flutter frequencies, the simulation and neurophysiological results show that the final firing rate depends significantly primarily on the sign of the difference and not on the magnitudes of the vibrotactile frequency values.

5.6.3 The probabilistic nature of the decision-making

The decision-making implemented by this attractor model is probabilistic. We now investigate this systematically.

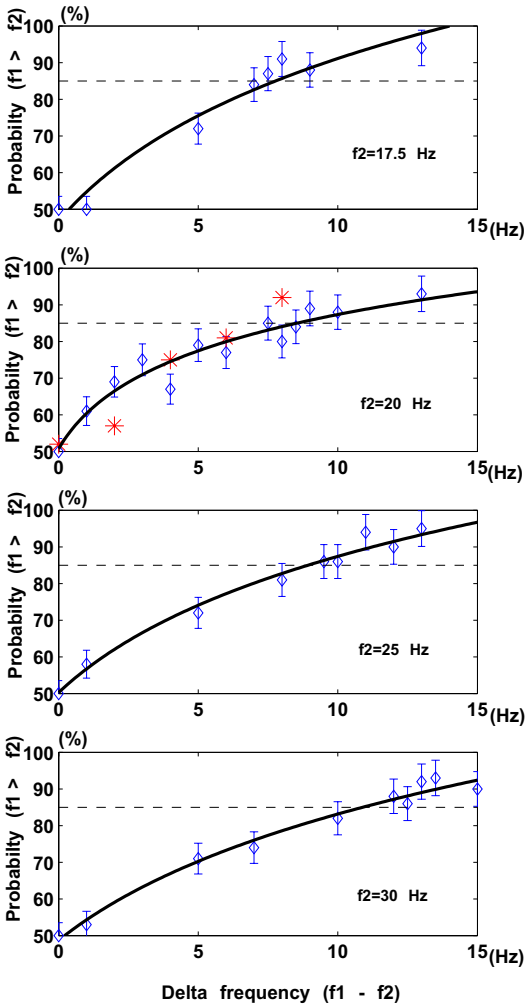

Fig. 5.7 Probability of correct discrimination ($\pm$ sd) as a function of the difference between the two presented vibrotactile frequencies to be compared. In the simulations, we assume that f1>f2 by a Δ-value (labelled 'Delta frequency (f1-f2)'), i.e. f1=f2+Δ. The points correspond to the trial averaged spiking simulations. The line interpolates the points with a logarithmic function. The horizontal dashed line represents the threshold of correct classification for a performance of 85% correct discrimination. The second panel down includes actual neuronal data (indicated by *) described by Romo and Salinas (2003) for the f2=20 Hz condition. (After Deco and Rolls 2006.)

Figure 5.7 shows the probability of correct discrimination as a function of the difference between the two presented vibrotactile frequencies to be compared. We assume that f1>f2 by a Δ-value, i.e. f1=f2+Δ. (In Fig. 5.7 this value is called 'Delta frequency (f1-f2)'.) Each diamond-point in the Figure corresponds to the result calculated by averaging 200 trials of the full spiking simulations. The lines were calculated by fitting the points with a logarithmic function. A correct classification occurs when during the 500 ms comparison period, the network evolves to a 'single-state' attractor that shows a high level of spiking activity (larger than 10 Hz) for the population (f1>f2), and simultaneously a low level of spiking activity for the population (f1<f2) (at the level of the spontaneous activity). Figure 5.7 shows that the

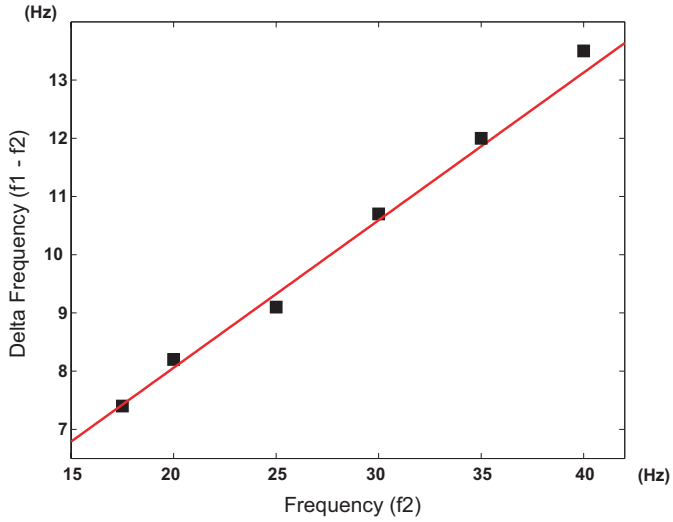

Fig. 5.8 Weber's law for the vibrotactile discrimination task. The critical discrimination Δ-value ('difference-threshold') is shown corresponding to an 85% correct performance level as a function of the base frequency f2. The 'difference-threshold' increases linearly as a function of the base frequency. (After Deco and Rolls 2006.)

decision-making is probabilistic, and that the probability of a correct discrimination increases as Δf, the difference between the two stimuli being compared, increases. When Δf is 0, the network performs at chance, and its choices are 50% correct. The second panel of Fig. 5.7 shows a good fit between the actual neuronal data described by Romo and Salinas (2003) for the f2=20 Hz condition (indicated by *), and the results obtained with the model (Deco and Rolls 2006).

One can observe from the different panels in Fig. 5.7 corresponding to different base vibrotactile frequencies f2, that to reach a threshold of correct classification of for example 85% correct (horizontal dashed line in Fig. 5.7), the difference between f1 and f2 must become larger as the base frequency f2 increases.

5.6.4 Probabilistic decision-making and Weber's law

Figure 5.8 plots the critical discrimination Δ-value corresponding to an 85% correct performance level (the 'difference threshold') as a function of the base frequency f2. The 'difference threshold' increases linearly as a function of the base frequency, that is, $\Delta \mathrm{f}/\mathrm{f}$ is a constant. This corresponds to Weber's law for the vibrotactile discrimination task. (Weber's law is often expressed as $\Delta I/I$ is a constant, where I stands for stimulus intensity, and ΔI for the smallest difference of intensity that can just be discriminated, sometimes called the just noticeable difference. In the case simulated, the stimuli were vibrotactile frequencies, hence the use of f to denote the frequencies of the stimuli.)

The analysis shown in Figs. 5.7 and 5.8 suggests that Weber's law, and consequently the ability to discriminate two stimuli, is encoded in the probability of performing a transition to the correct final attractor. To reinforce this hypothesis, Fig. 5.9 shows that Weber's law is not encoded in the firing rates of the VPC decision-making neurons that were modelled. Deco and Rolls (2006) simulated again a situation corresponding to the cases where f1 is larger than f2 by a Δ-value, and therefore the network will perform correctly when the dynamics perform a transition to the final single attractor corresponding to high activity in the population (f1>f2)

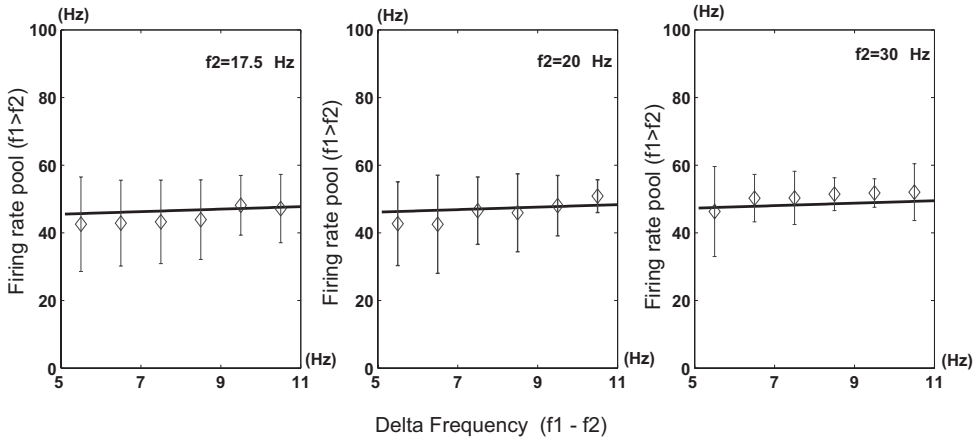

Fig. 5.9 Final stationary spiking rate of the (f1<f2) population (after transients) as a function of the difference between the two vibrotactile frequencies to be compared, for the cases where this population correctly won the competition and the network has performed a transition to the proper attractor. The plots show the results obtained with the spiking simulations. (The diamond-points correspond to the average values over 200 trials, and the error bars to the standard deviation.) The lines correspond to the mean-field calculations. The firing rate of the population encoding the result of the comparison is relatively independent of f2 and of the difference between f1 and f2, for the range of parameter values shown. (After Deco and Rolls 2006.)

and low activity in the population (f1<f2). Figure 5.9 plots for three different frequencies f2, the firing rate of the (f1<f2) population (for the cases where this population correctly won the competition and the network has performed a transition to the proper attractor) as a function of the difference between the two vibrotactile frequencies to be compared. The plots show the results obtained with the spiking simulations. The diamond-points correspond to the average values over 200 trials, and the error bars to the standard deviation. The lines show the results of the mean-field calculations. A good agreement between the spiking simulations and the corresponding mean-field results is observed. The most interesting observation is the fact that the firing rate of the population (f1>f2) in the correct attractor, for different base frequencies f2 and for different differences between f1 and f2 (Δf), is practically constant. Thus the firing rate of the population encoding the result of the comparison does not encode Weber's law. What happens is that the attractor dynamics makes a binary choice. The decision is reflected in a high firing rate of one population, and a low firing rate of the other population. These high and low firing rates thus reflect the binary choice, and depend rather little on Δf. Thus even for an equal value of f1 and f2, the attractor network makes a binary choice, with the probability of which choice is made reflecting Δf, but the firing rate that is reached as a result of the choice reflecting Δf very little. (In fact a small increase in firing rate is evident in Fig. 5.9 as Δf increases, and we show in Chapter 6 that this is related to decision confidence.)

The model also gives further insights into the mechanisms by which Weber's law is implemented. We hypothesized that because Δf/f is practically constant in the model, the difference of frequencies Δf required to push the single attractor network towards an attractor basin might increase with f because as f increases, shunting (divisive) inhibition produced by inhibitory feedback inputs (from the inhibitory interneurons) might act divisively on the pyramidal cells in the attractor network to shunt the excitatory inputs f1 and f2. In more detail, as the base frequency f increases, more excitation will be provided to the network by the inputs λ_1 and λ_2, this will tend to increase the firing rates of the pyramidal cells which will in turn provide a larger excitatory input to the inhibitory neurons. This will tend to make the

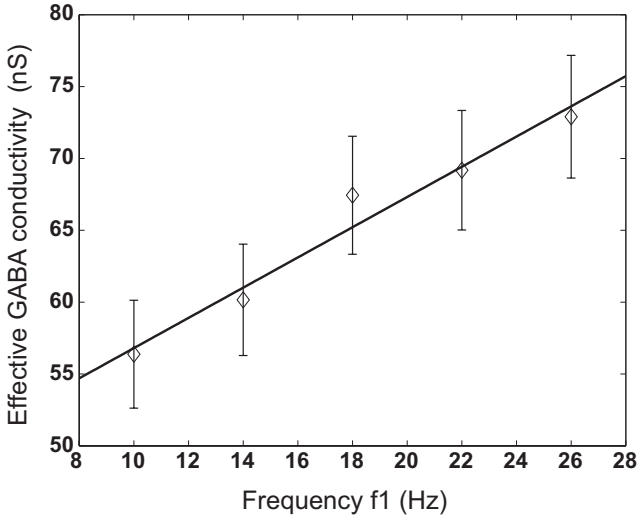

Fig. 5.10 The conductance in nS (mean $\pm$ sd) produced by the GABA inputs to the pyramidal cells as a function of the base frequency f1. The effective conductance produced through the GABA synapses (i.e. $I_{\mathrm{GABA}}/(V - V_{\mathrm{I}})$) was averaged over the time window in which the stimuli were presented in one of the excitatory neuron pools, when the base frequency was f1, and f2-f1 was set to 8 Hz. (After Deco and Rolls 2006.)

inhibitory neurons fire faster, and their GABA synapses onto the pyramidal cells will be more active. Because these GABA synapses open chloride channels and act with a driving potential $V_I = -70$ mV which is relatively close to the membrane potential (which will be in the range $V_L = -70$ mV to $V_{\mathrm{thr}} = -50$ mV), a large part of the GABA synaptic input to the pyramidal cells will tend to shunt, that is to act divisively upon, the excitatory inputs to the pyramidal cells from the vibrotactile biasing inputs λ_1 and λ_2. To compensate for this current shunting effect, f1 and f2 are likely to need to increase in proportion to the base frequency f in order to maintain the efficacy of their biasing effect. To assess this hypothesis, we measured the change in conductance produced by the GABA inputs as a function of the base frequency. Figure 5.10 shows that the conductance increases linearly with the base frequency (as does the firing rate of the GABA neurons, not illustrated). The shunting effect does appear therefore to be dividing the excitatory inputs to the pyramidal cells in the linear way as a function of f that we hypothesized.

Deco and Rolls (2006) therefore proposed that Weber's law is implemented by shunting effects acting on pyramidal cells that are produced by inhibitory neuron inputs which increase linearly as the base frequency increases, so that the difference of frequencies Δf required to push the network reliably into one of its decision attractors must increase in proportion to the base frequency. We checked the excitatory inputs to the pyramidal cells (for which $V_E = 0$ mV), and found that their conductances were much smaller (in the order of 5 nS for the AMPA and 15 nS for the NMDA receptors) than those produced by the GABA receptors, so that it is the GABA-induced conductance changes that dominate, and that produce the shunting inhibition.

Further properties of the attractor network model of decision-making that enable it to implement Weber's law include stability of the spontaneous firing rate condition, even when the decision cues are applied, so that the system depends on the noise to escape this starting state, as described in Section 5.8.

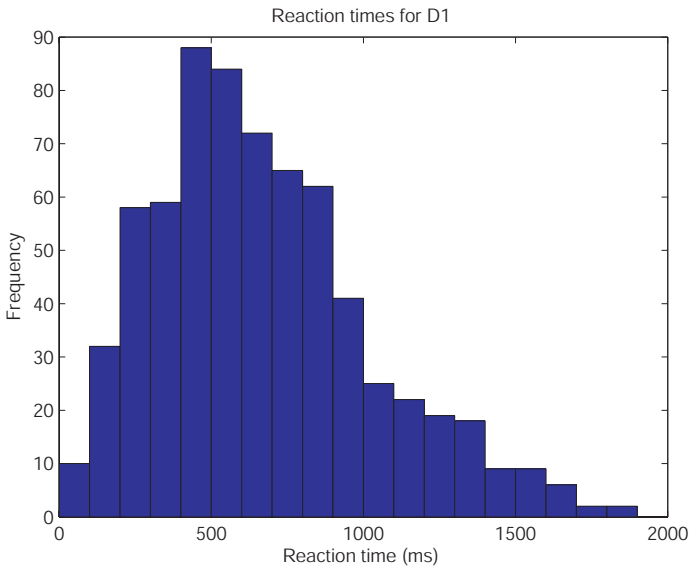

Fig. 5.11 Reaction time distribution for the decision-making attractor network described in Chapter 6. The difference between the two stimuli was relatively small (ΔI=16 Hz, though sufficient to produce 81% correct choices). The criterion for the reaction time was the time after application of both stimuli at which the firing rate of the neurons in the correct attractor became 25 Hz or more greater than that of the neurons in the incorrect attractor, and remained in that state for the remainder of the trial.

5.6.5 Reaction times

Because of the noise-driven stochastic nature of the decision-making, the reaction times even for one set of parameters vary from trial to trial. An example of the distribution of reaction times of the attractor network are shown in Fig. 5.11. This distribution is for a case when the difference between the two stimuli is relatively small (ΔI = 16 Hz for the simulations described in Chapter 6, though sufficient to produce 81% correct choices). The considerable variability of the reaction times across trials, and the long tail of the probability distribution, provide a very useful basis for understanding the variability from trial to trial that is evident in human choice reaction times (Welford 1980). Indeed, human studies commonly report skewed reaction time distributions with a long right tail (Luce 1986, Ratcliff and Rouder 1998, Ratcliff, Zandt and McKoon 1999, Usher and McClelland 2001, Marti et al. 2008).

The reaction times of this model of decision-making are faster when the discrimination is easy, in this case when the difference between f2 and f1 is large. This is analysed in Fig. 5.12. We calculated for a fixed f2=25 Hz and different f1>f2 (from 1 Hz to 13 Hz), the probability of correct or incorrect discrimination, and the corresponding reaction time. The reaction time was the time that the winning population ((f1>f2) for the correct cases, and (f1<f2) for the incorrect cases) took to cross a threshold of a firing rate of 20 Hz. (We considered the averaged reaction time over 200 trials.) Figure 5.12A plots the relation between the reaction time and the probability of correct classification. The larger the probability of correct classification, the faster is the decision-making, which is reasonable and consistent with the decision-making literature. Figure 5.12B plots the reaction time of an incorrect decision, as a function of the probability of performing a misclassification. The behaviour is now the converse, in that a low probability of incorrect discrimination implies also shorter reaction times. This means that the reaction time of a correct or incorrect classification is relatively similar, and therefore

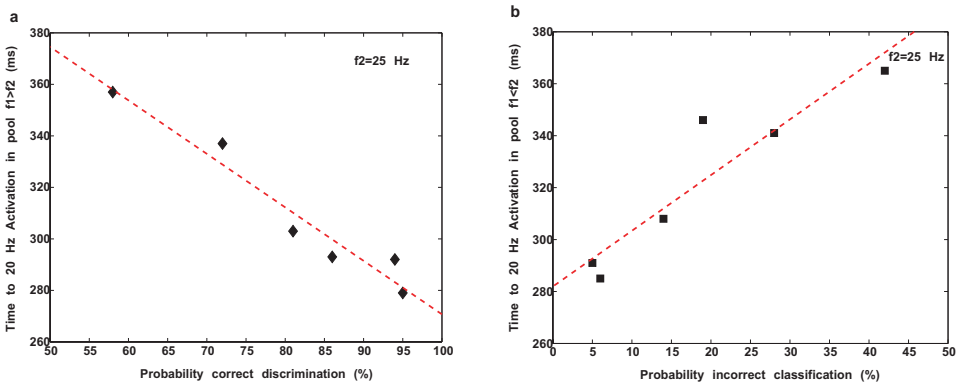

Fig. 5.12 Reaction time as a function of the probability of being correct or incorrect. Results are shown for a fixed f2=25 Hz and different f1>f2 (from 1 Hz to 13 Hz). The reaction time was the time that the winning population ((f1>f2) for the correct cases, and (f1<f2) for the incorrect cases) took to cross a threshold of 20 Hz. (The reaction time was averaged across 200 trials.) (a) Reaction time as a function of the probability of correct classification. The larger the probability the faster is the decision-making. (b) Reaction time of an incorrect decision, as a function of the probability of performing a misclassification. (After Deco and Rolls 2006.)

the dependence on the probability of a correct or incorrect classification is inverted.

Further analyses of reaction times as a function of the easiness of the decision are shown in Figs. 6.2, 6.3b, and 6.11.

5.6.6 Finite-size noise effects

The results described earlier indicate that the probabilistic settling of the system is related to the finite size noise effects of the spiking dynamics of the individual neurons with their Poisson-like spike trains in a network of limited size. The concept here is that the smaller the network, the greater will be the statistical fluctuations (i.e. the noise) caused by the random spiking times of the individual neurons in the network. In an infinitely large system, the statistical fluctuations would be smoothed out, and reach zero. To investigate this further, and to show what sizes of network are in practice influenced by these finite-size related statistical fluctuations, an important issue when considering the operation of real neuronal networks in the brain, Deco and Rolls (2006) simulated networks with different numbers of neurons, N. The noise due to the finite size effects is expected to increase as the network becomes smaller, and indeed to be proportional to $1/\sqrt{N}$. We show in Fig. 5.13 the effects of altering N on the operation of the network, where $N = N_{\mathrm{E}} + N_{\mathrm{I}}$, and $N_{\mathrm{E}} : N_{\mathrm{I}}$, was held at 4:1 as in the simulations shown earlier. The simulations were for f1=30 Hz and f2=22 Hz. Figure 5.13 shows overall that when N is larger than approximately 1,000, the network shows the expected settling to the (f1 > f2) attractor state on a proportion of occasions that is in the range 85–93%, increasing only a little as the number of neurons reaches 4,000 (top panel). The settling remains probabilistic, as shown by the standard deviations in the probability that the (f1 > f2) attractor state will be reached (top panel). When N is less than approximately 1,000, the finite size noise effects become very marked, as shown by the fact that the network reaches the correct attractor state (f1>f2) much less frequently, and in that the time for a decision to be reached can be premature and fast, as the large fluctuations in the stochastic noise can cause the system to reach the criterion [in this case of a firing rate of 20 Hz in the pool (f1>f2)] too quickly.

The overall conclusion of the results shown in Fig. 5.13 is that the size of the network,

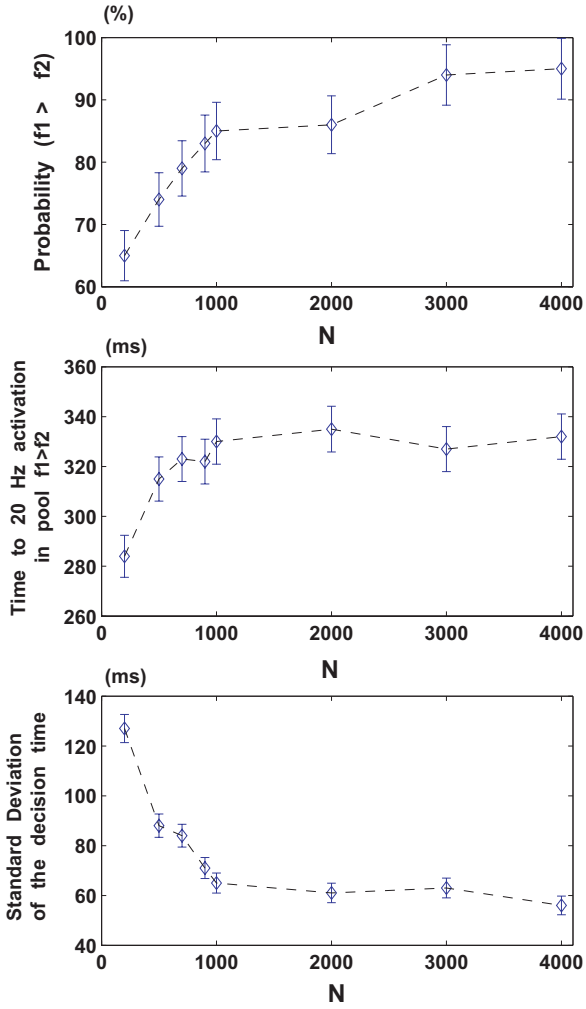

Fig. 5.13 The effects of altering N, the number of neurons in the network, on the operation of the decision-making network. The simulations were for f1=30 Hz and f2=22 Hz. The top panel shows the probability that the network will settle into the correct (f1>f2) attractor state. The mean $\pm$ the standard deviation is shown. The middle panel shows the time for a decision to be reached, that is for the system to reach a criterion of a firing rate of 20 Hz in the pool (f1>f2). The mean $\pm$ the standard deviation of the sampled mean is shown. The bottom panel shows the standard deviation of the reaction time. (After Deco and Rolls 2006.)

N, does influence the probabilistic settling of the network to the decision state. None of these probabilistic attractor and decision-related settling effects would of course be found in a mean-field or purely rate simulation, without spiking activity. The size of N in the brain is likely to be greater than 1,000 (and probably in the neocortex in the range 4,000–12,000) (see Table 1.1). With diluted connectivity, the relevant parameter is the number of connections per neuron (Amit and Brunel 1997).

It will be of interest to investigate further this scaling as a function of the number of neurons in a population with a firing rate distribution that is close to what is found in the brain, namely close to exponential (Franco, Rolls, Aggelopoulos and Jerez 2007) (see Section

1.11.3.2), as the firing rate distribution, and the sparseness of the representation, will influence the noise in the network. The results shown in Fig. 5.13 are for a sparseness of 0.1, and a binary firing rate probability distribution (i.e. one in which all the neurons in the winning attractor have the same average high firing rate, and the other neurons in the network have a low firing rate). Initial results of E.T.Rolls and T.Webb suggest that there is more noise with a biologically plausible approximately exponential firing rate distribution of the type illustrated in Fig. 1.14. The underlying cause may be the large noise contributed by the few neurons with rather high firing rates, given that for a Poisson distribution the variance increases with (and is equal to) the mean rate.

5.7 Properties of this model of decision-making

Key properties of this biased attractor model of decision-making are now described.

5.7.1 Comparison with other models of decision-making

In the attractor network model of decision-making described in this book and elsewhere (Wang 2002, Deco and Rolls 2006, Wang 2008, Deco, Rolls and Romo 2009b, Rolls, Grabenhorst and Deco 2010b, Rolls, Grabenhorst and Deco 2010c), the decisions are taken probabilistically because of the finite size noise due to spiking activity in the integrate-and-fire dynamical network, with the probability that a particular decision is made depending on the biasing inputs provided by the sensory stimuli f1 and f2.

The model described here is different in a number of ways from accumulator or counter models which may include a noise term and which undergo a random walk in real time, which is a diffusion process (Ratcliff, Zandt and McKoon 1999, Carpenter and Williams 1995) (see further Wang (2002), Wang (2008), and Usher and McClelland (2001)). First, in accumulator models, a mechanism for computing the difference between the stimuli is not described, whereas in the current model this is achieved, and scaled by f, by the feedback inhibition included in the attractor network.

Second, in the current attractor network model the decision corresponds to high firing rates in one of the attractors, and there is no arbitrary threshold that must be reached.

Third, the noise in the current model is not arbitrary, but is accounted for by finite size noise effects of the spiking dynamics of the individual neurons with their Poisson-like spike trains in a system of limited size.

Fourth, because the attractor network has recurrent connections, the way in which it settles into a final attractor state (and thus the decision process) can naturally take place over quite a long time, as information gradually and stochastically builds up due to the positive feedback in the recurrent network, the weights in the network, and the biasing inputs, as shown in Figs. 5.5 and 5.6.

Fifth, the recurrent attractor network model produces longer response times in error trials than in correct trials (Wong and Wang 2006, Rolls et al. 2010c) (see Chapter 6), consistent with experimental findings (Roitman and Shadlen 2002). Longer reaction times in error trials can be realized in the diffusion model only with the additional assumption that the starting point varies stochastically from trial to trial (Ratcliff and Rouder 1998).

Sixth,the diffusion model never reaches a steady state and predicts that performance can potentially improve indefinitely with a longer duration of stimulus processing, e.g. by raising the decision bound. In the recurrent attractor network model, ramping activity eventually stops as an attractor state is reached (Fig. 6.2). Consequently performance plateaus at sufficiently long stimulus-processing times (Wang 2002, Wang 2008).

Seventh, the attractor network model has been shown to be able to subtract negative signals as well as add positive evidence about choice alternatives, but the influence of newly arriving inputs diminishes over time, as the network converges towards one of the attractor states representing the alternative choices (Wang 2002, Wang 2008). This is consistent with experimental evidence (Wong, Huk, Shadlen and Wang 2007). This violation of time-shift invariance cannot be accounted for by the inclusion of a leak in the linear accumulator model. In fact, in contrast to the recurrent attractor network model, the linear leaky competing accumulator model, which takes into account a leakage of integration and assumes competitive inhibition between accumulators selective for choice alternatives (Usher and McClelland 2001), actually predicts that later, not earlier, signals influence more the ultimate decision, because an earlier pulse is gradually 'forgotten' due to the leak and does not affect significantly the decision that occurs much later (Wong et al. 2007).

The approach described here offers therefore a new, alternative, approach to this type of linear diffusion model, in the sense that the new attractor network model is nonlinear (due to the positive feedback in the attractor network), and is derived from and consistent with the underlying neurophysiological experimental data. The model thus differs from traditional linear diffusion models of decision-making used to account for example for human reaction time data (Luce 1986, Ratcliff and Rouder 1998, Ratcliff et al. 1999, Usher and McClelland 2001). The non-linear diffusion process that is a property of the attractor network model is analyzed further in Section A.6.1.

The model of decision-making described here is also different to a model suggested by Sugrue, Corrado and Newsome (2005) in which it is suggested that the probabilistic relative value of each action directly dictates the instantaneous probability of choosing each action on the current trial. The present model shows how probabilistic decisions could be taken depending on the two biasing inputs (λ_1 and λ_2 in Fig. 5.3, which could be equal) to a biased competition attractor network subject to statistical fluctuations related to finite size noise in the dynamics of the integrate-and-fire network.

5.7.2 Integration of evidence by the attractor network, escaping time, and reaction times

An interesting aspect of the model is that the recurrent connectivity, and the relatively long time constant of the NMDA receptors (Wang 2002), may together enable the attractor network to accumulate evidence over a long time period of several hundred milliseconds. Important aspects of the functionality of attractor networks are that they can accumulate and maintain information.

A more detailed analysis suggests that there are two scenarios that are needed to understand the time course of decision-making (Marti, Deco, Mattia, Gigante and Del Giudice 2008).

First, in the scenario investigated by Wang (2002), the spontaneous state is unstable when the decision cues are applied. The network, initially in the spontaneous state, is driven to a competition regime by an increase of the external input (that is, upon stimulus presentation) that destabilizes the initial state. The decision process can then be seen as the relaxation from an unstable stationary state towards either of the two stable decision states (see Fig. 5.14 (right)). When the system is completely symmetric, i.e. when there is no bias in the external inputs that favours one choice over the other, this destabilization occurs because the system undergoes a pitchfork bifurcation for sufficiently high inputs. The time spent by the system to evolve from the initial state to either of the two decision states is determined by the actual stochastic trajectory of the system in the phase space. In particular, the transition time increases significantly when the system wanders in the vicinity of the saddle that appears

when the spontaneous state becomes unstable. Reaction times in the order of hundreds of ms may be produced in this way, and are strongly influenced by the long time constants of the NMDA receptors (Wang 2002). The transition can be further slowed down by setting the external input slightly above the bifurcation value. This tuning can be exploited to obtain realistic decision times.

Second, there is a scenario in which the stimuli do not destabilize the spontaneous state, but rather increase the probability for a noise-driven transition from a stable spontaneous state to one of the decision states (see Fig. 5.14 (centre) and Section 5.8). Due to the presence of finite-size noise in the system there is a nonzero probability that this transition occurs and hence a finite mean transition rate between the spontaneous and the decision states. It has been shown that in this scenario mean decision times tend to the Van't Hoff-Arrhenius exponential dependence on the amplitude of noise in the limit of infinitely large networks. As a consequence, in this limit, mean decision times increase exponentially with the size of the network (Marti, Deco, Mattia, Gigante and Del Giudice 2008). Further, the decision events become Poissonian in the limit of vanishing noise, leading to an exponential distribution of decision times. For small noise a decrease in the mean input to the network leads to an increase of the positive skewness of decision-time distributions.

These results suggest that noise-driven decision models as in this second scenario provide an alternative dynamical mechanism for the variability and wide range of decision times observed, which span from a few hundred milliseconds to more than one second (Marti et al. 2008). In this scenario (see Section 5.8), there is an escaping time from the spontaneous firing state. In this time the information can be thought of as accumulating in the sense that the stochastic noise may slowly drive the firing rates in a diffusion-like way such than an energy barrier is jumped over and escaped (see Fig. 2.5 on page 73 and Fig. 2.6). In this situation, a landscape can be specified by a combination of the synaptic weights and external decision-related input evidence that biases the firing rates of the decision attractors, as described in Section 2.3. The noise introduced into the network by for example the random neuronal spiking can be conceptualized as influencing the way that the system flows across this fixed landscape shaped by the synaptic weights, and by the external inputs if they remain on during operation of the network.

The model for the second scenario, with a stable spontaneous state even when the decision cues are being applied, makes specific predictions about reaction times. One, in relation to $\Delta f / f$, is shown in Fig. 5.12. Further predictions about reaction times are shown in Figs. 6.3 and 6.11, and a fuller analysis showing that there will be a gamma-like distribution with an exponential tail of long reaction times in the reaction time distribution with this second scenario is provided by Marti et al. (2008).

5.7.3 Distributed decision-making

Although the model described here is effectively a single attractor network, we note that the network need not be localized to one brain region. Long-range connections between cortical areas enable networks in different brain areas to interact in the way needed to implement a single attractor network. The requirement is that the synapses between the neurons in any one pool be set up by Hebb-like associative synaptic modification, and this is likely to be a property of connectivity between areas (using forward and backprojections, see Section 1.9), as well as within areas (Rolls and Treves 1998, Rolls and Deco 2002). In this sense, the decision could be thought of as distributed across different brain areas. Consistent with this, Romo and colleagues have found neurons related to vibrotactile decisions not only in the ventral premotor cortex (VPC), but in a number of connected brain areas including the medial premotor cortex, as described in Section 5.3.

In order to achieve the desired probabilistic settling behaviour, the network we describe must not have very high inhibition, and, related to this, may sometimes not settle into one of its attractor states. In a forced choice task in which a decision must be reached on every trial, a possible solution is to have a second decision-making network, with parameters adjusted so that it will settle into one of its states (chosen at chance) even if a preceding network in the decision-making chain has not settled. This could be an additional reason for having a series of networks in different brain regions involved in the decision-making process.

In any case, we believe that there are decision-making networks, that is, networks that can reach a categorical state, in many cortical areas, each specializing in taking a decision about the information represented in that region (see Section 8.12). In this situation, a behavioural decision may reflect the operation of a number of partly separate, and partly sequential, decision-making processes.

5.7.4 Weber's law

Deco and Rolls (2006) showed that with this attractor-based model of decision-making, the relevant parameters for the decision to be made to a criterion of a given per cent correct about whether f1 is different from f2 by the network are found not to be the absolute value of f1 or f2, but the difference between them scaled by their absolute value. If the difference between the two stimuli at which they can be discriminated Δf = f1-f2, then it is found that Δf increases linearly as a function of the base frequency f2, which is Weber's law. The results show that Weber's law does not depend on the final firing rates of neurons in the attractor, but instead reflects the nature of the probabilistic settling into a decision-related attractor, which depends on the statistical fluctuations in the network, the synaptic connectivity, and the difference between the bias input frequencies f1 and f2 scaled by the baseline input f2.

Weber's law is usually formulated as $\Delta f / (f_0 + f)$ = a constant, where f_0 allows the bottom part of the curve to asymptote at f_0. In vision, f_0 is sometimes referred to as 'dark light'. The result is that there is a part of the curve where Δf is linearly related to f, and the curve of Δf vs f need not go through the origin. This corresponds to the data shown in Fig. 5.8.

An analysis of the non-stationary evolution of the dynamics of the network model, performed by explicit full spiking simulations, shows that Weber's law is implemented in the probability of transition from the initial spontaneous firing state to one of the two possible attractor states. In this decision-making paradigm, the firing rates of neurons in the VPC encode the outcome of the comparison and therefore the decision and motor response, but not how strong the stimuli are, i.e. what Weber called 'sensation' (as described for example in a detection task by de Lafuente and Romo (2005)). The probability of obtaining a specific decision, i.e. of detecting a just noticeable difference, is encoded in the stochastic dynamics of the network. More specifically, the origin of the fluctuations that will drive the transitions towards particular decisions depends on the connectivity between the different populations, on the size of the populations, and on the Poisson-like spike trains of the individual neurons in the system. In other words, the neural code for the outcome of the decision is reflected in the high rate of one of the populations of neurons, but whether the rate of a particular population becomes high is probabilistic. This means that an essential part of how the decision process is encoded is contained in the *synapses*, in the *finite size* of the network, and in the Poisson-like firing of individual neurons in the network.

The statistical fluctuations in the network are due to the finite size noise, which approximates to the square root of the (firing rate / number of neurons in the population) (see Mattia and Del Giudice (2002)), as shown in Fig. 5.13. This is the first time we know when the implementation of a psychophysical law is not the firing rate of the neurons, nor the spike timing, nor is single neuron based, but instead is based on the synaptic connectivity of the

network and on statistical fluctuations due to the spiking activity in the network.

The way in which the system settles (i.e. the probability of reaching one attractor state vs the other from the initial spontaneous state, and the time it takes) depends on factors that include the distortion of the attractor landscapes produced by the biasing inputs λ_1 and λ_2 which will influence both the shapes and the depth of the attractor basins, and the finite size noise effects. Of particular importance in relation to Weber's law is likely to be that when λ_1 and λ_2 increase, the increased firing of the neurons in the two attractors results in more activity of the inhibitory feedback neurons, which then produce effectively divisive inhibition on the principal cells of the attractor network. This is reflected in the conductance change produced by the GABA inputs to the pyramidal cells shown in Fig. 5.10. The inhibitory feedback is mainly divisive because the GABA-activated channels operate primarily as a current shunt, and do not produce much hyperpolarization, given that V_I is relatively close to the membrane potential. After the division implemented by the feedback inhibition, the differential bias required to push the network reliably into one of the attractors must then be larger, and effectively the driving force ($\lambda_1 - \lambda_2$ or $\Delta\lambda$) must get larger in proportion to the inhibition. As the inhibition is proportional to λ, this produces the result that $\Delta\lambda/\lambda$ is approximately a constant. We thus propose that Weber's law, $\Delta I/I$ is a constant, is implemented in part by shunting effects acting on pyramidal cells that are produced by inhibitory neuron inputs which increase linearly as the baseline input I increases, so that the difference of intensities ΔI required to push the network reliably into one of its attractors must increase in proportion to the base input I. Another part of the mechanism appears to be the noise-driven stochastic process of jumping out of a stable state of spontaneous firing in what we describe as a multistable operating regime, as described in Section 5.8.

We emphasize that this account (Deco and Rolls 2006) of Weber's law is intended to be a general account, and is not restricted to the particular data set or brain system against which the development of the model described here was validated.

A prediction of the model is that Weber's law for frequency discrimination could be implemented not by the firing rate of a given population of neurons (which reflects just the discrete decision taken), but by the probability that a particular population will be activated, which depends on Δf / f. This prediction could be tested by a trial-by-trial analysis of the neurophysiological data in which the firing of neurons at different base frequencies f and for different Δf is measured, to investigate whether the type of result shown in Fig. 5.7, and thereby in Fig. 5.8, which are derived from the model, are also found in the neuronal data from the experiments, and this would also usefully confirm that Weber's law holds at the neuronal level in this particular vibrotactile task with a delay. In particular, although Romo, Hernandez and Zainos (2004) in their figure 5 show that choice probability and neuronal activity increases as a function of f2-f1, we predict that neurons should follow the functions shown in Figs. 5.7 and 5.8 for different values of Δf and f, and it would be of interest to test this prediction.

We note that Weber's law holds in most though not all discrimination situations, and to the extent that Weber's law does generally hold, the model described here provides a computational neuroscience-based account of how it arises. This is the first time we know when the implementation of a psychophysical law is not the firing rate of the neurons, nor the spike timing, nor is single neuron based, but instead is based on the synaptic connectivity of the network and on statistical fluctuations due to the spiking activity in the network.

5.7.5 Unifying principles

In summary, we now have a model for how attractor networks could operate by biased competition to implement probabilistic decision-making. This type of model could operate in very

many brain areas, as described in Section 8.12. This model of decision-making is part of the larger conceptual issue of how memory networks retrieve information. In this case the short-term memory property of the attractor network helps the network to integrate information over time to reach the decision. Although retrieval of information from attractor networks has been intensively studied using *inter alia* approaches from theoretical physics (see Section 1.10, Amit (1989), and Rolls and Treves (1998)), the way in which the retrieval is probabilistic when an integrate-and-fire implementation with spiking neurons is considered has been less studied. The approaches of Brunel and Wang (2001) and Deco and Rolls (2006) do open up the issue of the probabilistic operation of memory networks with spiking dynamics.

We may raise the conceptually important issue of why the operation of what is effectively memory retrieval is probabilistic. Part of the answer is shown in Fig. 5.13, in which it is seen that even when a fully connected recurrent attractor network has 4,000 neurons, the operation of the network is still probabilistic. In this network of 4,000 neurons, there were 3,600 excitatory neurons and 360 neurons represented each pattern or decision (that is, the sparseness a was 0.1). The firing rates of the neurons corresponded to those found in VPC, with rates above 20 spikes/s considered to be a criterion for the attractor state being reached, and 40–50 spikes/s being typical when fully in the attractor state (see Fig. 5.9). Under these conditions, the probabilistic spiking of the excitatory (pyramidal) cells is what introduces noise into the network. (Deco and Rolls (2006) showed that it is this noise in the recurrent collateral firing, rather than external noise due to variability in the inputs, which makes the major contribution to the probabilistic behaviour of the network.) Thus, once the firing in the recurrent collaterals is spike implemented by integrate-and-fire neurons, the probabilistic behaviour seems inevitable, even up to quite large attractor network sizes.

To investigate the 'decision-related' switching of these systems, it may be important to use a firing rate distribution of the type found in the brain, in which few neurons have high rates, more neurons have intermediate rates, and many neurons have low rates (see Section 1.11.3). It is also important to model correctly the proportion of the current that is being passed through the NMDA receptors (which are voltage-dependent), as these receptors have a long time-constant, which will tend to smooth out short-term statistical fluctuations caused by the stochastic firing of the neurons (cf. Wang (1999)), and this will affect the statistics of the probabilistic switching of the network. This can only be done by modelling integrate-and-fire networks with the firing rates and the firing rate distributions found in a cortical area.

Reasons why the brain is inherently noisy are described in Section 2.4.

Applications of this model of decision-making are described in Chapter 8.

5.8 A multistable system with noise

In the situation illustrated in Figs. 2.5 and 5.1, there is multistability, in that the spontaneous state and a large number of high firing rate persistent states are stable. More generally, and depending on the network parameters including the strengths of the inputs, a number of different scenarios can occur. These are illustrated in Fig. 5.14. Let us consider the activity of a given neuronal population while inputs are being applied.

In Fig. 5.14 (left) we see a situation in which only the spontaneous state S is stable. This might occur if the external inputs λ_1 and λ_2 are weak.

On the right we have a situation in which our neuronal population is either in a high firing rate stable state C2, or in a low firing rate state C1 because another population is firing fast and inhibiting our neuronal population. There is no stable spontaneous state.

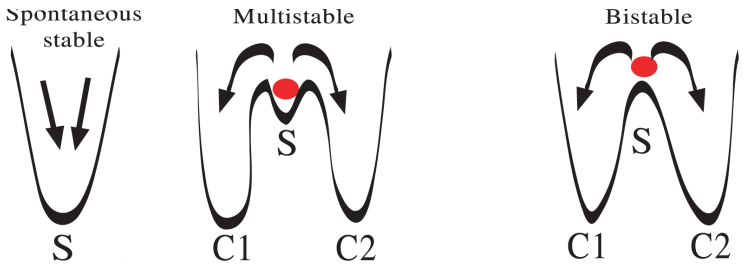

Fig. 5.14 Computational principles underlying the different dynamical regimes of the decision-making attractor network (see text). The x-axis represents the neuronal activity of one of the populations (ν_i) and the landscape represents an energy landscape regulating the evolution of the system. S is a stable state of spontaneous activity, C2 is a high firing rate state of this neuronal population corresponding to the choice implemented by this population, and C1 is a low firing rate state present when the other population wins the competition.

In the middle of Fig. 5.14 we see a situation in which our population may be either in C1, or in C2, or in a spontaneous state of firing S when no population has won the competition. We emphasize that this can be a scenario even when the decision cues λ_1 and λ_2 are being applied during the decision-making period. We refer to this system as a multistable system.

The differences between these scenarios are of interest in relation to how noise influences the decision-making. In the scenario shown in the middle of Fig. 5.14 we see that there are three stable states when the inputs λ_1 and λ_2 are being applied, and that it is the stochastic noise that influences whether the system jumps from the initial spontaneous state to a high firing rate state in which one of the decision-state populations fires fast, producing either C2 if our population wins, or C1 if our population loses. The statistical properties of the noise (including its amplitude and frequency spectrum), and the shape of the different basins in the energy landscape, influence whether a decision will be taken, the time when it will be taken, and which high firing rate decision attractor wins. In contrast, in the scenario shown in Fig. 5.14 (right) the energy landscape when the stimuli are being applied is such that there is no stable spontaneous state, so the system moves to one of the high firing rate decision attractors without requiring noise. In this case, the noise, and the shape of the energy landscape, influence which high firing rate decision state attractor will win.

In an exploration of the neural network described in this Chapter that models vibrotactile flutter frequency discrimination in humans, Deco, Scarano and Soto-Faraco (2007b) found that the neurodynamical mechanisms and computational principles underlying the decision-making processes in this perceptual discrimination task are consistent with a fluctuation - driven scenario in a multistable regime of the type illustrated in Fig. 5.14 (middle) as incorporated in the model of Deco and Rolls (2006). Deco et al. (2007b) used a mean-field analysis with a system of nonlinear coupled differential equations of the Wilson-Cowan type (Renart et al. 2003, La Camera et al. 2004) to describe the evolution of the average firing rate of each population, and added a noise fluctuation term to drive the transitions.

Fig. 5.14 represents schematically the different ways in which the network could operate, depending on the value of λ, the input, as shown by a mean-field analysis. The x-axis represents the neuronal activity of one of the populations (ν_i) and the landscape represents an energy landscape regulating the evolution of the system. The energy landscape reflects the synaptic values and the effects of the incoming sensory information (the λ_1 and λ_2 values in Fig. 5.1). For values of $\lambda < \lambda_{c1}$=20 Hz (Fig. 5.14, left panel), only the spontaneous state is stable, and no decision states appear (for the unbiased case, i.e. when $\lambda_1 = \lambda_2$). For increasing $\Delta\lambda$ (biased case), one decision state (corresponding to the choice where the increased

value of $\lambda + \Delta\lambda$ is applied) emerges, attracting the dynamics towards this decision state. For values of λ between λ_{c1}=20, λ_{c2}=40 Hz (Fig. 5.14, middle panel) there is a region of multi-stability, with the spontaneous state and each of the two decision states all stable. In this λ interval, the fluctuations are responsible for driving the system from the initial stable spontaneous state to one of the two decision states corresponding to the two possible response choices. (C2 is a high firing rate of the population shown in Fig. 5.14 middle that corresponds to the decision implemented by that population of neurons.) Thus, in this scenario, fluctuations play a crucial role in the computation of decision-making. It is only in this region that Weber's law is found to apply to the network, as described in this Chapter. For values of $\lambda > \lambda_{c2}$=40 Hz (Fig. 5.14, right panel) a region of bistability is found where the initial spontaneous state is unstable, and only the two decision states are stable. In this regime, the spontaneous state destabilizes, so that the dynamics rapidly evolves towards one of the two decision states, resembling therefore a pure diffusion integrating the relative evidence for one choice over another. The fact that the network can implement Weber's law, and does so only in a range of values for λ in which the network operates as a fluctuation-driven multistable system, provides further evidence to support the hypothesis that decision-making is implemented by a multistable fluctuation-driven attractor system, where there are in the unbiased case stable states for the spontaneous firing state, and for each of the decisions.

With respect to the multistable scenario (in the middle of Fig. 5.14), the attractor network acts as a short-term memory that can accumulate evidence over time, and usually gradually though stochastically the firing rates of the two groups of neurons corresponding to the two choices diverge, one set of neurons stochastically increasing their firing rate, and the other set being inhibited by the first set of neurons. A situation like this is probably occurring in Fig. 5.5, and is analyzed in more detail by Marti et al. (2008). We can illustrate this point by the middle landscape in Fig. 5.14, and note that the accumulation of evidence corresponds to the position in the space indicated by the ball moving noisily in a direction towards for example C2, but not yet jumping over the energy barrier into the C2 attractor.

It is important to stress that the effect of the noise is particularly relevant in the multistable regime (middle of Fig. 5.14), because the fluctuations are the driving force that allow the system to escape the decision barriers around the stable spontaneous state. In the multistable scenario, the choices are associated with stable attractors, and the starting condition is also given by a stable spontaneous state. To make a decision, the system has to escape the stable spontaneous state towards one of the choice-attractors. (This is related to the so called 'Kramers' escape problem' (Kramers 1940)). On the other hand, in the bistable regime (right of Fig. 5.14), the so called 'ballistic' regime), the noise is of course relevant as the basis of the diffusion process, but it is not the main driving force. This is because in the bistable scenario the spontaneous state is not a stable state, and therefore with or without noise, the system will necessarily evolve to one or the other decision-attractor just because of the neurodynamical flow (Deco and Marti 2007a, Deco and Marti 2007b).

Another interesting aspect of the model is that the recurrent connectivity, and the relatively long time constant of the NMDA receptors (Wang 2002), together enable the attractor network to accumulate evidence over a long time period of several hundred milliseconds. This is an important aspect of the functionality of attractor networks. Nevertheless, long reaction times can also be obtained without NMDA receptors, using the alternative scenario of multistability (Marti, Deco, Mattia, Gigante and Del Giudice 2008). In this case, the level of noise is the main variable that drives the escape from the stable spontaneous firing rate state, and low levels of noise can produce long reaction times.